

Adaptive Regulation: Instrument Choice for Policy Learning over Time

Lori S. Benneer and Jonathan B. Wiener*

Abstract: Regulation is often viewed as a one-time, go/no-go decision, based on ex ante forecasts of uncertain future benefits and costs. But the world changes, and static past policies can become mismatches with emerging new science, technology, and social conditions. For example, a single decision could yield suboptimal errors by overregulating a net beneficial technology, under-regulating a net harmful technology, or spurring side effects or other unintended consequences. There is thus interest in moving from static to adaptive regulation, replacing the one-time yes/no decision with a series of multiple partial sequential decisions informed by monitoring and review – i.e., designing regulation to incorporate learning over time. Moving to a more adaptive approach poses both pros and cons, including the gains of improving net benefits, reducing policy errors, and overcoming political impasse, but also the costs of data collection, decision analysis, and policy instability. And there can be different institutional mechanisms for adaptive regulation. Hence we need an instrument choice framework for adaptive regulation. We distinguish between unplanned adaptive (e.g. crisis response, ad hoc retrospective review) and planned adaptive (e.g. periodic review, adaptive licensing), and within planned adaptive regulation we further distinguish between discretionary and automatic mechanisms. We develop a decision tree for deciding among adaptive regulatory instruments, considering criteria such as the frequency of review, the scope of impacts, and implications for the political economy of regulation.

Keywords:

JEL Codes: K23, K32, L51, L98, O38, Q48

* Lori S. Benneer is the Juli Plant Grainger Associate Professor of Energy Economics and Policy at the Nicholas School of the Environment (primary), Sanford School of Public Policy and Department of Economics and Duke University. She is the Associate Director for Educational Programs at the Duke University Energy Initiative. Jonathan B. Wiener is the Perkins Professor of Law, Environmental Policy and Public Policy at the Duke Law School (primary), the Sanford School of Public Policy and the Nicholas School of the Environment. Benneer and Wiener are co-directors of the Rethinking Regulation program at the Kenan Institute for Ethics at Duke University. We thank Ashley Shan and Christine Gerbode for helpful research assistance.

1. Introduction

Consider a spectrum from perfectly static to perfectly adaptive policy: at the perfectly static end, a policy once promulgated never changes, while at the perfectly adaptive end, a policy updates instantaneously in response to new information. Neither extreme reflects the real world, but they bracket the range of options one can envision for incorporating learning over time into regulatory policy. In some cases regulation can effectively be a one-time decision, difficult to revise except over long time periods, rather than a series of multiple sequential provisional decisions that can be revised more quickly. It is perhaps telling that under US administrative law, a promulgated agency regulation is called a “final” rule. (The imaginary type of perfectly adaptive policy recalls the notion of perfectly efficient markets, which instantly incorporate new information, whereas real markets adapt through actual human cognition and behavior, see Lo 2017.)

Critics worry that modern regulation is too static, adopting rigid requirements that will become mismatches as the world changes through evolving science, technology, conditions and values (e.g. Marchant et al. 2011, 2013). This can occur even where initial policy is based on the best efforts at ex ante regulatory impact assessments, because uncertain forecasts of benefits and costs can later prove inaccurate (e.g. Aldy 2014, Morgenstern 2015). Policy may tend to be static for several reasons, including political gridlock, institutional inertia, and behavioral heuristics such as single action bias and loss aversion. Numerous regulatory scholars have argued that the regulatory process needs to become more adaptive (e.g. Farber 1994; Botkin 1996; Freeman & Farber 2005; McCray Oye & Petersen 2010; Marchant et al. 2011; Marchant et al. 2013; Wiener 2002, 2004, 2017; Wiener & Richman 2010; Wiener & Ribeiro 2016; Pidot 2015; Craig et al. 2017; DeCaro et al. 2017; Cosens et al. 2018; Carlson and Burtraw,

forthcoming). Recent research finds that pressure for adaptive revision already generates a degree of dynamism through informal agency actions, but that these revisions may favor particular interest groups and do not always adhere to transparency procedures such as notice and comment rulemaking (Wagner et al. 2017).

Moving toward more adaptive regulation offers several potential advantages. Static policies in the face of change can yield accumulating over-regulation of net beneficial technology, under-regulation of net harmful technology, and adverse side effects and other unintended consequences. Such change can be exogenous, but it can also be in response to the policy, yielding unintended ancillary impacts (Graham & Wiener 1995). Thus, more adaptive regulation could offer reduced policy errors and countervailing risks, improved net benefits, and potentially more ambit for technological innovation. These outcomes could also help to lessen public frustration over outdated and obsolete policies. In addition, more adaptive regulation could potentially overcome some political impasses, easing the initial adoption of policies by promising further chances to assess and revisit the policy later (including more opportunities for public input on policy revisions).

The case in favor of more adaptive regulation is especially strong when we can anticipate that we will learn a great deal after the initial adoption of a policy. Advancing technologies -- for example, highly autonomous vehicles, genetic engineering and editing, personalized medicine, and nanotechnology -- present enormous promise to improve quality of life while simultaneously presenting potential risks that are uncertain and changing over time. Regulation of these emerging technologies still follows the classic U.S. approach to regulation which typically invests a great deal of effort and time in a single rulemaking process in which problems are identified, alternative solutions are evaluated, a solution is proposed and open for public

comment, and then a final rule is issued. Yet emerging technologies are inherently dynamic. The fast pace of change creates significant uncertainties for regulators. The risk profile—how much risk the technology presents and to whom—is changing quickly, as is the set of possible alternative regulatory strategies, which mean that the net benefits profile from regulation of the technology may also be changing rapidly. These technologies particularly warrant an approach to regulation that adjusts and adapts as information about the technology changes.

It may be that today’s technologies demand more adaptive and flexible regulatory approaches than in the past. If so, that could be because these new technologies are disrupting the traditional model of innovation. Older “hardware” based technologies tended to be produced in centralized facilities, and hit the market fully formed, with issues and kinks worked out before release. New “software” based or distributed technologies are often released early, use data from early adopters to refine and edit the technology, and issue frequent updates and improvements. These technologies generally embrace a more adaptive approach to innovation that now demands a more adaptive approach to regulation. Meanwhile, it may be that more generally, beyond the cases of current emerging technologies, other changes (such as in scientific understanding, economics, ecology, demographics, and social values) could also warrant adaptive regulatory approaches that learn and update over time. Many environmental and economic systems that have traditionally been regulated in a static manner may actually be dynamic, warranting more adaptive approaches to their regulation as well (e.g. Botkin 1990; Botkin 1996; Wiener 1995, 1996; Adler 2016).

On the other hand, moving toward more adaptive regulation may pose its own drawbacks. The costs of adaptive regulation include the cost of data collection, which may incur public expenditures and may also be burdensome on those required to provide the data (Sunstein 2019).

(The costs of some kinds of data collection may be falling as information technology expands through sensors, drones, and the internet of things, but those in turn may raise privacy concerns.) The analysis and decisions to review and revise a policy may be costly, to the agency itself in terms of scarce staff time and the opportunity costs of other policies not addressed (which is one reason that busy agencies may resist devoting time to retrospective reviews), as well as to society if the updated policy is more costly. And, despite the need for regulatory change to adapt to these changes in the world, there is also value in the predictability of legal rules. Roscoe Pound recognized both sides of this tradeoff when he counseled that “Law must be stable, and yet it cannot stand still” (Pound 1923). Policy instability can impose costs, especially if unanticipated (as we discuss below regarding “unplanned” adaptive regulation). Policy stickiness may be warranted in some cases to avoid post-adoption erosion of public protections by powerful interest groups (Lazarus 2009, Coffee 2012). The perception of policy instability may erode public confidence in the rules (although so may obsolete rules that fail to keep up with change). If adaptive regulation looks too facile, or is seen as favoring interest groups rather than promoting social well-being, it may call into question the credibility of the government’s commitment to stick to the initial rules, thereby undermining compliance. Thus, it may not always be worthwhile to move regulation toward more adaptive approaches. In some cases the gains will not justify the costs.

Amidst these pros and cons, the spectrum from static to adaptive is not just one step: there are several different ways that regulatory institutions could fashion adaptive mechanisms to incorporate learning over time. Each instrument for adaptive regulation may pose a different mix of pros and cons. In general, policy learning can occur through several paths, including through learning from observed variation in policies across jurisdictions or across agencies

(Listokin 2008; Vogel 2017), purposefully designed experiments (Greenstone 2009 Stavins 1998, Schmalensee & Stavins 2013), crisis surprise and response (Balleisen et al. 2017b), retrospective review of past policies (Farber 1994, Coglianese 2013, Aldy 2014, Sunstein 2014, Morgenstern 2015, Wiener & Ribeiro 2016), sunset clauses that terminate policies pending review (Gubler 2014), adaptive licensing of new products to limited subpopulations (Eichler et al. 2012), and planned ongoing monitoring and periodic review (McCray Oye & Petersen 2010), among others.

Hence we need not just a debate over static vs. adaptive, but a framework for instrument choice among the different approaches to adaptive regulation. We need to evaluate these instrument choices in terms of the pros and cons just noted. Key elements of evaluating the move from static toward adaptive will involve the frequency of review (posing the tradeoff of more rapid updating vs. greater instability), the scope of impacts to be assessed (posing the tradeoff of full social impact assessment including unintended ancillary impacts, vs. the costs of greater data collection and analysis), the decision maker (posing the tradeoff of automatic vs. discretionary – rapid and reliable, but based on a limited scope of impacts, vs. slower and more deliberate, but assessing a broader scope of impacts), and the implications for the political economy of regulation (such as easing initial adoption vs. undermining subsequent compliance).

The literature on adaptive regulation is hindered by a lack of a common lexicon. Different scholars use the term “adaptive regulation” differently and different definitions lead to different notion of what policy instruments are “adaptive.” Moreover, there are references to “flexible regulation” and “responsive regulation” which sound similar to “adaptive regulation” and are occasionally used interchangeable while occasionally used to mean something else entirely.

This paper seeks to help chart a framework of instrument choice for adaptive regulation. First, we define adaptive regulation in broad terms, and contrast unplanned adaptive with planned adaptive regulation. Second, we subdivide planned adaptive regulation into discretionary adaptive and automated adaptive regulation, and give examples of both types, as well as examine some of the key issues involved when choosing either approach. Third, we develop a decision tree for determining whether adaptive regulation should be used, and if so, whether it should be discretionary or automated.

2. Toward a General Definition of Adaptive Regulation

A standard dictionary definition of “adapt” is “to make something suitable for a new use or purpose; modify.” Thus, a concept of adaptive regulation could be regulation that can be modified or changed in response to (and to be suitable for) new situations. In trying to develop a definition (or definitions) of adaptive regulation, we can also look to other disciplines that have wrestled with this concept of adaptiveness. In ecology and environmental law, for example, and especially regarding ecosystem management, the notion of “adaptive management” is a structured process of iterative learning from prior management decisions in order to respond to change and improve policy outcomes (Holling 1978; Craig and Ruhl 2014; Cosens et al. 2018; Craig et al. 2017; DeCaro et al. 2017).

As a first effort at defining adaptive regulation, we could simply transform the definition of adaptive management to fit the regulatory process. This would lead to a definition of adaptive regulation as a structured process of iterative learning from the impacts of prior regulatory decisions. We expand this definition to include the outcome of this learning, namely that policy changes over time. We also modify the definition to include what determines the outcome/change, namely data collection and analysis. Hence, our definition of adaptive

regulation is: “a structured regulatory process that enables learning and modification of policy over time via adjustments informed by data collection and analysis” Or: “laws built to learn.” Clearly critical in this definition is that what is adapting over time is the policy itself, and that these changes are not ad hoc, but rather part of a systematic review that draws on evidence and analysis.

The concept of “planned adaptive regulation” (PAR) (McCray Oye & Petersen 2010; IRGC 2017) adds the idea that the regulatory system is planned to adapt, i.e. that its adaptive quality is designed from the outset. PAR contemplates regulation that incorporates learning – that is built to learn and update over time. (While this distinction is helpful, it may also be redundant if the term adaptive itself connotes a process by which changes can be made based on new information. But as we discuss below, there could also be “unplanned adaptive” regulation, in which regulation adapts in response to a surprise, without prior planned data collection and review.)

Having settled on a definition does not necessarily eliminate confusion. There are several other terms used in the regulatory literature that sound quite similar to adaptive regulation but may mean quite different things. These terms include: flexible regulation, responsive regulation, self-regulation, and co-regulation. Understanding how adaptive regulation differs from these other typologies of regulation requires an examination of the components of a regulation.

Every regulation could be considered to consist of four main components: the regulator, the target, the commands, and the consequences (Bennear and Coglianese 2013), as depicted in Figure 1. A canonical command-and-control regulation would consist of a government agency as the regulator, a specified target to reduce some social harm, a command to use a technology or

process to meet the target, and a set of inspections and penalties to ensure compliance with the command (though not necessarily evaluating performance in terms of reducing the social harm). Over time, the literature on regulation has focused on myriad ways to move away from this standard command-and-control framework by broadening or relaxing different aspects of the regulatory quadrumvirate. The terms used for all of these various broadenings of the framework sound similar and can be confusing.

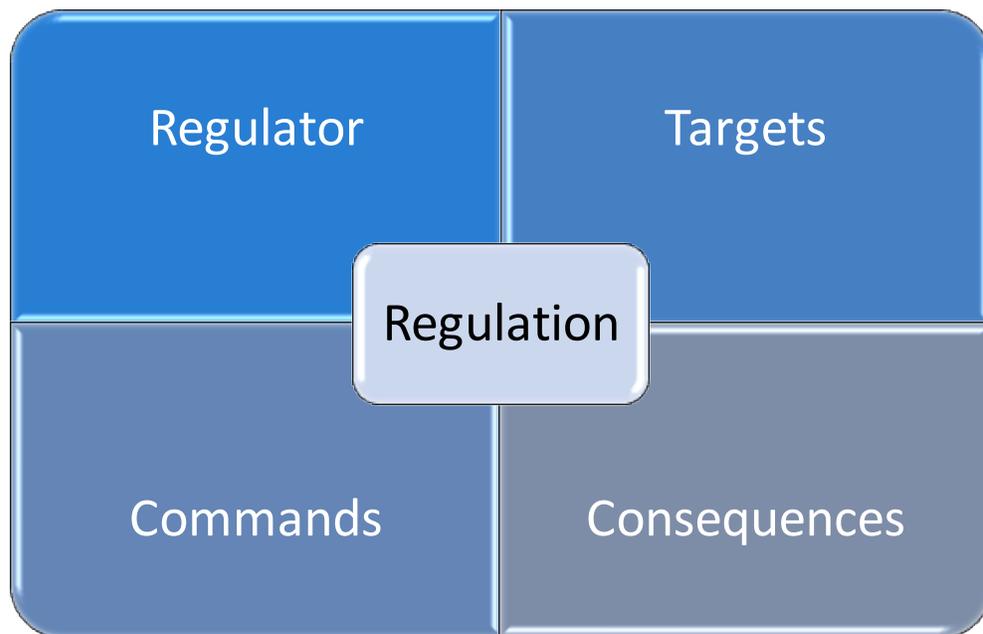


Figure 1: The regulatory quadrumvirate

The terms “self-regulation” and “co-regulation” have been used to describe situations in which industry serves as both the regulator and the regulatee or in which industry and government develop the regulation collaboratively. Hence, these terms focus on the upper-left

(regulator) and upper-right (target) quadrant in the regulatory matrix (Maxwell, Lyon, and Hackett 1998; King and Lenox 2000; Benneer and Coglianese 2013).

The term “flexible” regulation has been used to describe situations in which a regulation enables regulated entities to change their behavior to achieve the target and/or in the face of new information (Carlson and Burtraw forthcoming). A policy such as a performance standard is flexible regulation in that the regulator specifies a target but leaves the regulated entities the flexibility to determine how best to meet the target. Hence, the term “flexible regulation” is focused on the commands, which are minimal other than requiring that the regulated entities achieve the target. While this distinction is very helpful, it has not always been maintained in the literature. In our own work (Benneer and Coglianese 2013) we have previously argued that flexibility can apply to all four of the components of regulation. Nonetheless, it seems clarifying and helpful to develop different terms for different types of flexibility and reserve the term “flexible” regulation for regulations that have flexible or non-prescriptive commands. Cap and trade systems, which (like performance standards) authorize regulated entities to have flexibility in how best to meet the target (the method of abatement), also add further flexibility for regulated entities to re-allocate the location of abatement under the aggregate cap by trading allowances across emitters (and potentially to re-allocate abatement over time, via banking and borrowing). These quadrants of regulator, target, command and consequences, with associated degrees of flexibility, also find their analogs in the typology of regulatory policy instruments in Wiener and Richman (2010), including conduct instruments (command and control technology/design requirements), quantity instruments (including performance standards, cap and trade, and property rules), price instruments (including taxes and liability rules), and information instruments (disclosure rules). Much of the literature on “flexible” regulation has

emphasized “how” flexibility in the methods of abatement (e.g. via performance standards rather than technology requirements), and “where” flexibility in the spatial location of abatement (e.g. via tradable allowances or taxes rather than uniform standards); it has also introduced some mechanisms to enable “when” flexibility in the timing of abatement (e.g. via banking and borrowing allowances) (see the review in Pizer & Prest 2018). These types of flexibility have examined the optimal initial choice of instrument in the face of uncertainty (e.g. the classic paper by Weitzman 1974) and then on learning and updating adjustments by the regulated entities in response to the initially set choice of policy instrument (e.g. Tarui & Polasky 2005). Adaptive regulation, by contrast, emphasizes the temporal dimension through learning and updating adjustments by the regulator. (Below we note that some examples of adaptive regulation over time have been raised in the literature on flexible regulation, such as price ceilings and price floors in allowance trading markets. And recent work by Pizer & Prest 2018 finds that policy updating, in combination with banking and borrowing and other conditions, can yield superior outcomes.)

The concept of “responsive regulation” sounds similar to adaptive regulation, but has a different emphasis. Responsive regulation is focused on the nature of the relationships between regulators and stakeholders, including regulated entities (Ayres and Braithwaite 1994), particularly with respect to how to promote compliance with regulations. Responsive regulation promotes a two-way flow of communication between regulated entities and regulators to reward pro-societal behavior and promote compliance (and even over-compliance) with regulations. Hence responsive regulation is really focused on broadening the opportunities for stakeholder participation in the decision process.

We also distinguish adaptive regulation from “incrementalism” or “muddling through” (Lindblom 1959), to the extent that Lindblom urged that regulatory agencies should avoid “comprehensive” analysis in which “every important relevant factor is taken into account,” and instead take an incrementalist approach in which “analysis is drastically limited” and “important possible outcomes are neglected” (Lindblom 1959: 81). He added that the practice of “ignoring important possible consequences” might seem “a shocking shortcoming” but that it was superior to “futile attempts to achieve a comprehensiveness beyond human capacity” (Lindblom 1959: 85). He expressed confidence that other agencies would address the adverse side effects of each agency’s actions, in a “division of labor [in which] every important interest or value has its watchdog” (Lindblom 1959: 85-86). Yet the reality is often that government is fragmented, agencies exhibit bounded decision making, and affected populations lack effective voice, so that incrementalism fails to address the adverse side effects of regulatory actions (Graham & Wiener 1995; Wiener 1998; Stewart 2014). Hence Lindblom’s incrementalism may be adaptive, but with minimal scope of analysis of impacts, inviting serious countervailing risks. Another approach is labeled “experimentalism” by Sabel and co-authors, referring to regulation that sets broad performance objectives while empowering local regulatory officials to gather in-depth data on internal industry practices about complex technologies and to develop detailed knowledge to require best practices (Sabel & Simon 2011). But this approach is not the same as experimentation across policy types (Greenstone 2009). This “experimentalism” is one policy tool and can be compared to other policy tools in studies of observed variation or purposeful policy experimentation. Indeed, the policy type that Sabel and Simon juxtapose to their version of “experimentalism”—the use of economic incentive systems, such as taxes or tradable permits, which they label “minimalism” (although they need not be at all “minimalist”

in stringency, scope, or outcomes) – were themselves ambitious “experiments” (Stavins 1998, Schmalensee & Stavins 2013). In this sense, “experimentation” is broader than the term “experimentalism,” to the extent that the latter is one type of regulation whereas the broader term “experimentation” refers to conducting experiments to compare across multiple types of regulation.

Meanwhile, decentralization to local authorities is highlighted as a feature of adaptive regulation or adaptive governance by some authors (e.g. Brunner 2010, Cosens et al. 2018, Craig et al. 2017, Sabel & Simon 2011), but this spatial characteristic can be seen as a version of learning from variation across jurisdictions (Listokin 2008) and can be distinguished from the temporal dimension of learning and updating over time in adaptive regulation.

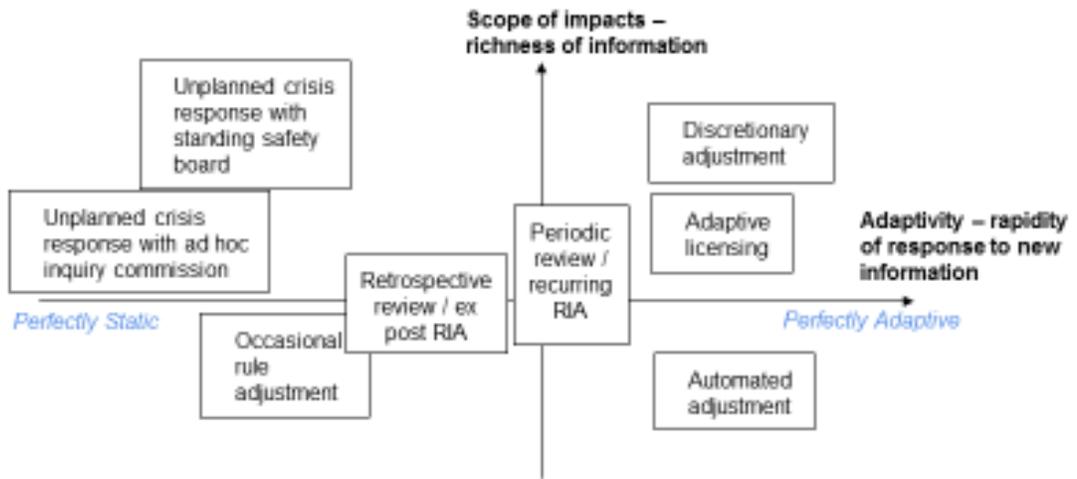
Within the category of adaptive regulation, there remains a spectrum of adaptivity (see McCray Oye & Petersen 2010; Wiener & Ribeiro 2016). Consider these points along this spectrum: (i) A one-time regulatory decision with no follow-up and no change could be deemed non-adaptive or static. (ii) A regulatory decision with a single subsequent follow-up evaluation (“look-back” or “retrospective review” or “ex post impact assessment” or “follow up evaluation”) begins to be somewhat adaptive (Coglianese 2013, Aldy 2014, Sunstein 2014), but the single follow-up may itself be ad hoc, and may be interposed only years after the initial regulation with no effort having been undertaken to monitor, collect data, and analyze impacts (consequences) over time (Wiener & Ribeiro 2016). (iii) A regulatory decision is more adaptive if it plans ahead to collect data and undertake a (single) subsequent follow-up evaluation (at a future point in time for retrospective/ex post review). Such a follow-up retrospective review can occur with or without a “sunset clause” – that is, the default option can be that the regulation

continues until it is modified (no sunset), or that it ends unless it is renewed (sunset) (Gubler 2014). (iv) Even more adaptive would be a planned series of ongoing monitoring, data collection, and analysis, with periodic evaluation of the consequences (recurring RIA) and iterative updating/revision (McCray Oye & Petersen 2010; IRGC 2017; Ribeiro 2018). Such ongoing evaluation with periodic iterative updating puts a focus on questions such as: how frequent the intervals for evaluation and revision should be (periodicity), which impacts or consequences should be monitored (scope), and which institutions should have the authority to undertake each task and adopt which types of revisions (power).

Another source of variation within adaptive regulation turns on how automatic the adaptations are. The examples above all suggest that a regulatory agency (or surrogate) engages in periodic review and revision of the regulation. That is, there is active participation of the regulator in the process of adaptive updating. However, there are also examples where regulations are initially established with thresholds or triggers, such that if those thresholds are reached or triggers are activated, the rule automatically changes in predefined ways. This type of regulation also fits our definition of adaptive regulation. We refer to regulation that requires direct action by a regulator to change the rule after the initial promulgation as Discretionary Adaptive Regulation. In contrast, Automated Adaptive Regulation establishes the rules and terms of adaptation in the initial rulemaking and the regulation changes over time automatically based on realized outcomes. (Pidot 2015 offers a related typology; his type of “contingent” policy is similar to our type of automated adaptive policy.) It is possible that an adaptive regulation could be a hybrid, with some discretionary options to update and some automated updating.

Figure 2 presents the spectrum from static to adaptive, with several (but not all) of the instruments arrayed along this spectrum. Notably, Figure 2 illustrates not only the degree of adaptivity (or rapidity of response to new information), shown on the horizontal axis, but also the scope of impacts assessed and hence the richness of information needed to make the adaptive decision, on the vertical axis.

Figure 2: Spectrum of Adaptivity and Scope



3. Typology of Adaptive Regulation

3.1. Unplanned Adaptive Regulation

We have spoken of traditional regulation as static, in the sense that a decision to adopt (or not) a statute and then an implementing rule is initially envisioned to endure indefinitely. Truly static regulation would remain rigid despite the world changing. (One example of which we are aware is Japan's national ambient air quality standards, which apparently remained the same from their adoption in the 1970s through 2009.) Yet in reality we do of course observe legal change over time, such as through legislative amendments or repeals, and through administrative agency revisions or rescissions of past rules. Much of these changes are “unplanned adaptive,” in the sense that the changes were not initially envisioned, and no plan was initially adopted to collect monitoring data and undertake periodic reviews. As noted above, a single retrospective review (follow up evaluation), undertaken ex post without an initial plan to collect and analyze data, could be a form of constrained unplanned adaptive regulation. A policy adopted with ex ante impact assessment and a plan to monitor data toward a single planned retrospective ex post evaluation could be seen as a partial but constrained planned adaptive regulation.

Moreover, some or many “unplanned adaptive” regulatory changes occur not due to retrospective follow up evaluation, but after surprises – crisis events that upend past static regulation (Percival 1998; Birkland 2006; Repetto 2006; D. Carpenter and Sin 2007; Wuthnow 2010). Examples include oil and chemical spills, nuclear power accidents, transportation accidents, workplace safety accidents, foodborne disease outbreaks, drinking water contamination, hazardous waste revelations, and more. Such “policy shocks” can dramatize the

need for regulatory change, reveal the arguable shortcomings of past regulation, spark contested framings of blame, and open windows of political opportunity for policy entrepreneurs to advance new approaches (Balleisen et al. 2017b).

To be sure, not every crisis spurs new regulation, and not every new regulation arises from a crisis (Kahn 2007). The point here is that when regulatory change does occur in response to a crisis or surprise, is typically “unplanned adaptive” in the sense that it adapts to changes in the world (such as a disaster, a scandal, new technology, new scientific understanding, or new social values), but without a prior plan to undertake monitoring and periodic review and analysis.

Regulatory change after a surprise or crisis can be needed, indeed overdue, to address neglected or festering problems. But where it lacks careful ongoing preparation and analysis, such unplanned adaptive regulatory change can be wrenching. Under the exigent circumstances in the wake of a crisis, underanalyzed regulatory change can make errors of diagnosis and remedy, and can yield new rules that are poorly designed, inadequately effective, excessively costly, and/or generate adverse ancillary impacts (countervailing risks). Unanticipated legal change can also impose losses (and gains) that actors have not insured against, spurring demands for compensation (Kaplou 1986, 1992, 2003) and political pressure for subsequent policy reversals.

The lack of institutions to enable learning from crises and design better regulatory approaches is a pitfall of unplanned adaptive approaches, giving rise to policy regret (Balleisen et al. 2017b). As we discuss below, creating a standing “safety board” or investigative body, ready to work whenever a surprise occurs, and thus accumulating institutional expertise over time on how to diagnose and respond, can be a significant improvement over unplanned crisis response via a commission of inquiry created anew from scratch (Balleisen et al. 2017a). The

lack of opportunity to learn is a particularly vexing problem of extreme risks that are so rare and so catastrophic (e.g. ending all life, or destroying civilization) that learning cannot be undertaken after the event – rendering adaptive learning insufficient, and putting a premium on precaution (Wiener 2016). Designing and preparing institutions for monitoring, analysis, and updating is essential to move from static to unplanned adaptive to planned adaptive regulation.

3.2. Planned Adaptive Regulation

The ambition of planned adaptive regulation (PAR) is to design regulatory systems that are built to learn. Key features include planned data collection, analysis, and periodic review, with opportunities to modify the initial policy (McCray Oye & Petersen 2010; IRGC 2017). Several key dimensions of PAR may be need to be calibrated to fit the issue being addressed, including:

- Periodicity: how frequently should periodic reviews be conducted?
- Scope: which variables should be monitored (including not only target impacts and costs but also unintended ancillary impacts, distributional impacts, etc.)?
- Comparison: to what counterfactual baseline, and to which other alternative policy options, should the actual policy be compared to measure its performance over time and the potential net benefits of updating? Can adaptive reviews help test and improve the accuracy of ex ante impact assessment forecasts?
- Predicate for policy change: Should updating be discretionary or automatic?
- Power and institutional roles: Who should collect data, analyze, undertake periodic reviews, and decide on policy changes?
- Politics: could moving from static to unplanned adaptive to PAR affect the political economy of regulation, including during the PAR process, and also ex ante (how interest

groups may change strategies when they anticipate changeable PAR rather than static enduring rules)?

- Meta-net benefits: The value of improved policy approaches from PAR needs to be compared to the information and decision costs of PAR. And thus, one must consider ex ante whether PAR should be applied to all issues, or only to selected key problems posing the highest anticipated net benefits of adaptivity. And then, one can envision a further layer of ex post or ongoing planned adaptive review of PAR itself, to see how well PAR is actually working in practice and make adjustments over time.

3.2.1. Discretionary Adaptive Regulatory Instruments

Discretionary adaptive regulatory instruments typically take the form of structured opportunities for analysis of regulatory performance by the regulator or an oversight or advisory body, coupled with an opportunity for regulatory revision at the discretion of the regulator. An example is the Federal Reserve Board's Open Markets Committee which regularly considers and sometimes adjusts interest rates.

Retrospective reviews of regulations (sometimes called retrospective regulatory impact analyses or ex post impact assessment) require agencies to "look back" at the effectiveness of a regulation and potentially also measure the realized benefits and costs of that regulation (Coglianese 2013; Aldy 2014; Sunstein 2014). This type of retrospective review can be a form of discretionary adaptive regulation, but only (or more so) if it is coupled with real opportunities to monitor, learn, and adjust the regulation being reviewed (Wiener & Ribeiro 2016; Ribeiro 2018). For example, Executive Order 13563, signed by President Barack Obama in 2011, required

Section 6(a). To facilitate the periodic review of existing significant regulations, agencies shall consider how best to promote retrospective analysis of rules that may be outmoded, ineffective, insufficient, or excessively burdensome, and to modify, streamline, expand, or repeal them in accordance with what has been learned. Such retrospective analyses, including supporting data, should be released online whenever possible.

The goal was to use retrospective review of significant regulations as a means of learning and reforming such regulations. However, the executive order was limited in its ability to truly foster adaptive regulation because it required the review but did not require subsequent adaptive action based on those reviews. Furthermore, it left agencies with significant discretion to determine which significant rules to review, using what data, compared to which baseline and alternatives, and with what periodicity (Aldy 2014; Morgenstern 2015; Wiener & Ribeiro 2016; Cropper Fraas & Morgenstern 2017). Other Presidents had also issued executive orders seeking retrospective review of existing rules (e.g. President Clinton's EO 12886; and President Carter's EO 12044), with limited results (see Wiener & Ribeiro 2016).

Another version of discretionary adaptive regulation is adaptive licensing (Eichler et al. 2012). This instrument is relevant to products that require licensing before they can be put on the market. For example, pharmaceutical drugs are traditionally evaluated for licensing based on a social risk/benefit evaluation, and then either approved or not – once, for everyone. A new approach is adaptive licensing, in which a new drug is initially approved only for a limited subpopulation (e.g. those for whom it may offer the greatest net benefits), and then this initial population is studied to observe efficacy and side effects in practice (going beyond smaller clinical trials), with careful analysis of genetic and other factors that may help predict health effects in the broader population. (This approach implies that the regulator has some way to constrain access to the product to just the approved subpopulation, which may not be the case

under all laws in all countries.) After this initial subpopulation, further licensing decisions may be made for other subpopulations or for all. Another possible application for this adaptive licensing approach may be automated vehicles, approving or mandating them for those who would most benefit from them (e.g. repeat drink drivers) and studying those results before approving them for the entire population. (This approach may face tradeoffs with the challenge of mixing fleets of automated and human drivers, and with the heightened public perception of initial accidents involving automated vehicles.) Adaptive licensing enables the regulator to learn over time and for different subpopulations, rather than having to make a once-and-for-all decision with more limited information.

One way to ensure that reviews not only occur but are used to adapt regulation is to require periodic review and revision. The US Clean Air Act requires periodic reviews of the national ambient air quality standards (NAAQS) for criteria air pollutants, to determine whether revisions to the rule are warranted. If the enabling legislation does not require such a periodic review and revision, the agency could decide to engage in that review anyway. However, the incentives for agencies to engage in meaningful review when it is not required are limited. Reviews are costly and potentially embarrassing to the regulator, and revisions can be politically heated. Unless the regulation is truly not fit for purpose, it is difficult for an agency to justify the expense (both monetary and political) of re-examining a rule that has been successfully promulgated and implemented. Furthermore, if the regulation has been fully implemented by industry, then an analysis that shows that the regulation was overly burdensome or not particularly effective, while potentially informative for future rule-makings, may not actually do much to reduce costs of the existing regulation which are already sunk.

The potential costs and benefits of retrospective review and regulatory revision also affect the desired periodicity of such a review. The US Clean Air Act section 109 requires a review of the NAAQS every 5 years; the new Lautenberg Chemical Safety Act adds to TSCA a provision in section 2625(l) for review every 5 years. But one might reasonably ask whether five years is the optimal time frame for such reviews. Ideally the frequency should be connected to the expected gain in learning – which depends on the rate at which change is occurring, new information is emerging, and net benefits are evolving. If the net benefits profile is changing more slowly and won't shift significantly for say 10 years, then a five-year periodicity is wasteful. Alternatively, if change is rapid and net benefits are shifting annually, a five-year periodicity may be too infrequent.

A key policy question when using discretionary adaptive regulation is about institutional power -- “who” should be the discretionary actor. One choice is to require or empower the regulatory body itself to review the data, conduct new analysis, and propose changes, if needed. The advantages of this approach are that the regulatory agencies are often best technically qualified and historically equipped to understand the topic area being regulated, to collect and understand the data, and to review the prior analysis. The downsides of asking the regulator to review regulations are that the agency may be very busy under pressure to address new issues, and there may be little incentive for the regulator to find flaws in its own regulatory regime, if only because such failures may be embarrassing and offer ammunition for political foes (Wiener & Ribeiro 2016).

A slight modification of this approach would be to have the agency conduct the adaptive analyses and propose changes, but have analyses and changes reviewed by an oversight body. The executive branch already has this type of oversight for benefit-cost analysis of major

regulations—significant new regulations must have their ex ante regulatory impact assessments reviewed by the Office of Information and Regulatory Affairs (OIRA) in the Office of Management and Budget (OMB). While current staffing levels at OIRA are not sufficient for it to take on the additional task of reviewing all retrospective reviews, analyses, and proposed changes, or ongoing adaptive reviews, this task could be well-suited to the expertise of that agency should resources be made available.

Alternatively, one could imagine having the analyses and proposals for changes conducted or at least reviewed by a third party. While there is not currently a model for this approach, one could imagine modifying existing institutions to handle this task. One existing model that might be modified to fit this role is the accident safety boards, such as the National Transportation Safety Board (NTSB) and the Chemical Safety Board (CSB). The best of the safety boards have developed deep expertise in investigation processes that can be used to investigate accidents regardless of the domain (Balleisen et al. 2017a). Other models for such third-party reviews might be at the National Academy of Sciences (NAS) or the General Accountability Office (GAO), or a newly created commission. One might imagine a regulatory review board that develops deep expertise in the process of regulatory review and adaptation that can be applied across policy domains. The advantages of this approach are that it better allows the review to be critical of the current regulatory approach, when appropriate; and that it allows consistent application of expertise and methodologies for review across different agencies and policy domains. Ensuring the appropriate scope of impacts, finding the optimal periodicity, comparing to counterfactual baselines, improving the accuracy of ex ante RIA, weighing decision costs and value of information, and other methodological issues may arise similarly in numerous retrospective or adaptive reviews. The disadvantages are that this organization is

unlikely to be able to develop substantive policy expertise in all policy domains, that it would require significant funding and organizational change to implement, and that such a third party body would be unlikely to have the authority to implement policy changes itself (as long as its reviews would be advisory to the regulatory agencies).

A key disadvantage to the use of discretionary adaptive regulation is that policy change must be consciously undertaken. In a polarized political climate, policy change often seems impossible. But even in the most cooperative and bi-partisan political environments, policy change can be politically difficult. Once a particular policy is implemented, its beneficiaries tend to build political coalitions to either continue the policy (especially in situations of loss aversion, and concentrated benefits with diffuse costs). Regardless of the technocratic merits or de-merits of a particular regulation, these political coalitions are likely to play a key role in any potential discretionary adaptation. (On the other hand, *ex ante*, the anticipation of the potential future adaptability of a regulation may facilitate initial agreement by actors who might otherwise resist any adoption of a static rule that did not meet their preferences. If so, adaptive regulation could enable the adoption of socially desirable policy changes that static regulation would not.)

3.2.2 Automated Adaptive Regulatory Instruments

By contrast, automated adaptive regulatory instruments do not require the regulator (or any other body) to actively make changes after the promulgation of the regulation. The idea of an automated adaptive regulatory instrument is that the conditions under which the rule will change are established as part of the initial rulemaking. This may come in the form of predetermined thresholds or triggers. If a particular threshold is reached, then the rule will change, but the exact nature of the change and the threshold itself are defined in the initial rule.

An example may help clarify what we mean by automated adaptive regulation. Consider a cap and trade program for emissions control. A simple cap and trade system, one where the cap is established and permit allocation is determined in the rulemaking, so that firms can subsequently trade permits among themselves, would not meet our definition of planned adaptive regulation. Such a regulation would be flexible (as described above), but not adaptive, because it has no process by which the cap itself can be changed in the face of new information on the marginal benefits or marginal costs of pollution.

However, researchers have proposed ways to make the cap change in response to new information and hence to transform the simple cap and trade system from a flexible regulation to one that is both flexible and adaptive. The use of price-floors (a price at which the government will buy excess permits) or price-ceilings (a price at which the government will sell additional permits) can actually change the goal/mandate of the program in an adaptive way (Murray, Newell, and Pizer 2009). For example, imagine a pollution permit trading system with both a price floor and price ceiling. If marginal pollution abatement costs are much lower than anticipated (perhaps because of advancing abatement technology), the price of permits falls. When the price hits the price floor, the government starts to buy any additional permits that are being sold at that price. This reduces the number of permits and increases the stringency of the regulation. This is exactly what one would want from a social-welfare maximization standpoint. If marginal costs are lower than anticipated (holding marginal benefits fixed), we should do more pollution abatement. Similarly, if the marginal abatement costs end up being much higher than anticipated (perhaps because of higher than expected costs for low-emitting technologies), the price of permits rises. If the price hits the price ceiling, the government sells additional permits at that price. This reduces the stringency of the regulation, which is again, exactly the desired

effect from a welfare-maximization standpoint. If the price floors and ceilings are determined in the initial rule, then the regulation is automatically adaptive. If the price floor or ceiling thresholds are reached, the cap is automatically changed.

The Waxman-Markey climate bill, which passed the House in 2009 and failed in the Senate in 2010, contained a version of a cap and trade system that was automatically adaptive. During the bill's formation, there were concerns about the uncertainty surrounding how much carbon abatement would cost. If abatement costs were higher than expected, some legislators wanted assurances that the cap could be loosened. However, other legislators, concerned about ensuring that emissions not be allowed to continue unchecked, wanted to ensure that there were firm limits on total carbon emissions. The compromise was to set the cap on carbon emissions but create a permit reserve equal to at most 10 percent of total permits. If carbon prices exceeded a target price (beginning in 2014 this target price would have been 160 percent of the 3-year rolling average of daily prices), permits could be released from the reserve up at that price. Since the reserve contained a fixed number of additional permits, this still guaranteed some emissions reductions relative to a business-as-usual scenario (Maniloff and Murray 2009).

The Waxman Markey bill also had a price floor, below which no additional permits would be sold. This price floor serves a similar purpose in allowing the overall cap to adjust based on realized marginal benefits and marginal costs. In Waxman Markey the price floor started at \$10 and increased each year by 5 percent plus inflation. If prices fall below the floor, no additional permits are offered at the auction, thus lowering the cap (Maniloff and Murray 2009). (Pizer & Prest 2018 examine another version of adaptive policy in a cap and trade system, with banking and borrowing and also policy updating over time.)

Other cap and trade programs have adopted similar adaptive measures. The Regional Greenhouse Gas Initiative (RGGI), which controls carbon emissions in the Northeastern United States, includes an auction price floor and price ceiling that triggers a permit reserve. Recent reviews of the RGGI program have led to consideration of more price steps between the floor and the ceiling, allowing the overall cap to adjust more frequently based on marginal costs and marginal benefits (Burtraw et al. 2017).

Automated adaptive regulation could go further, to be more truly and fully automated, through the use of algorithms, machine learning and artificial intelligence (AI) to revise policy in response to new information and a changing world (Coglianese & Lehr 2017). Such automated adaptive regulation could be more rapid in responding to multiple kinds of change, but may also raise questions of transparency, explainability, accountability, embedded bias, and other aspects (IRGC 2018). To be sure, these shortcomings must be compared to discretionary human decisions, which may also suffer from heuristics and biases. And new methods of computer science may assist in overseeing algorithm based decisions (Kroll et al. 2017).

3.3. Comparing Automated and Discretionary Approaches

There are advantages and disadvantages to both the discretionary and automated adaptive regulatory approaches. One approach does not strictly dominate the other. Two of the key tradeoffs between discretionary and automated approaches concern the speed of adaptation and the accuracy of adaptation.

In general, automated adaptive regulations result in more rapid adaptation to new information, since no additional analysis or rulemaking is required and all of the adaptation takes place in the context of the original rule. In contrast, discretionary adaptive regulatory

instruments require a detailed process of review and adjustment which could take years, by which time the process may need to begin again as the technologies have further changed resulting in additional changes to the net benefit profile. Equally important is the political cost of discretionary adaptive regulation. If discretionary adaptation requires a new notice-and-comment rulemaking process, the adaptation is then subject to political conflict and political costs that may not justify whatever improvements the adaptation would bring.

Discretionary adaptation does have some benefits over automated adaptation, however. The tradeoff of the speed in automated adaptive regulatory instruments may be a loss in accuracy or appropriateness of the adaptations. If all of the automated adaptation rules must be designated in the initial rulemaking, this requires good knowledge about the key parameters that will trigger future adaptation, at the time of the original rulemaking. For example, in the case of a price floor in a pollution permit system, new technologies that lower costs of carbon abatement may cause the price to fall below the floor. The pre-set price floor will help maintain the stringency of the cap, but ex post, it may turn out to be more efficient to raise the price floor itself (tighten the cap even further) and reduce pollution even further. That option is eliminated or restricted with an automated approach. (Perhaps the legislature could intervene to tighten the cap, in a version of discretionary adaptation by another institution, with its own decision costs and political conflicts.)

Also, if the side effects of a regulation (countervailing harms or co-benefits) are unexpectedly large, a broader scope of impact evaluation may be needed to learn and a discretionary adaptive process may be needed to revise the regulation (Graham & Wiener 1995; Wiener 2002; Wiener & Ribeiro 2016). A discretionary adaptive regulatory instrument can

better embrace the full scope of information and impacts that may be revealed on net benefits of technologies and policies.

A potential middle ground to reconcile the speed-accuracy tradeoff may be realized via a hybrid approach: if the regulation sets up an automated adaptive instrument, and also creates processes for discretionary adaptation of those automated instruments over time. The RGGI carbon market essentially has done this, by designating a price floor and ceiling for the auction, but also requiring periodic programmatic reviews that may change those market parameters (Burtraw et al. 2017). Perhaps algorithm and AI approaches to automated adaptive policy can also build in opportunities for recursive learning and adjustments, either by the program, by human discretionary oversight, or both.

Another way to reduce delays in a discretionary adaptation framework is to have the initial rulemaking (or legislation) establish a panel or body that is authorized to review and update the regulatory parameters without having to go through a new rulemaking process. For example, if the initial regulation establishes an adjustment board that is empowered to use data to adjust the some parameter of the regulation, then this board could evaluate both benefits and costs and change standards without the delay of additional notice-and-comment rulemaking. Such adjustments would raise questions about public input to important policy changes. And such adjustments will still be slower than the adjustment that comes from automated adaptive instruments, but the additional breadth of adaptiveness could warrant this added time, at least in some situations. An example of this regularized discretionary adaptive adjustment approach could be the Federal Reserve Board's Open Markets Committee, which can adjust interest rates fairly quickly; or the Montreal Protocol's processes for adding to its set of covered substances (such as CFCs, HCFCs, and HFCs).

Figure 3 below sketches a decision framework for regulators in assessing options for the predicate of planned adaptive regulation. The first node in the decision tree is whether adaptive regulation is likely to be beneficial at all. If the net benefits profile is relatively static over time, then the additional costs of adaptation are unlikely to be warranted and a non-adaptive approach makes the most sense.

If the net benefits profile is likely to change considerably over time, the next node in the decision tree is whether the regulator has enough information to establish thresholds for key uncertain parameters at the time of the initial rulemaking. If the answer is “yes,” then an automated adaptive approach is warranted. Clearly, this requires some professional judgement as there is always uncertainty and it is usually likely to be reduced over time. So, in essence, this decision node requires regulators to weigh how much uncertainty they have over key parameters now, how much this uncertainty will be resolved over time, the net benefits of the automatic adjustment mechanism, and the political costs of changing the regulation later.

This brings us to the third node of the decision tree which is about the political costs of changing the regulation. Some regulations are sufficiently controversial that the option to re-open the political dialogue for review in the future is extremely costly. If the political appetite for potential adaptations is there, then a discretionary adaptive approach may be best. If the political costs are viewed as high, then a more automated adaptive approach may make sense even if the key parameters for such adaptation are uncertain.

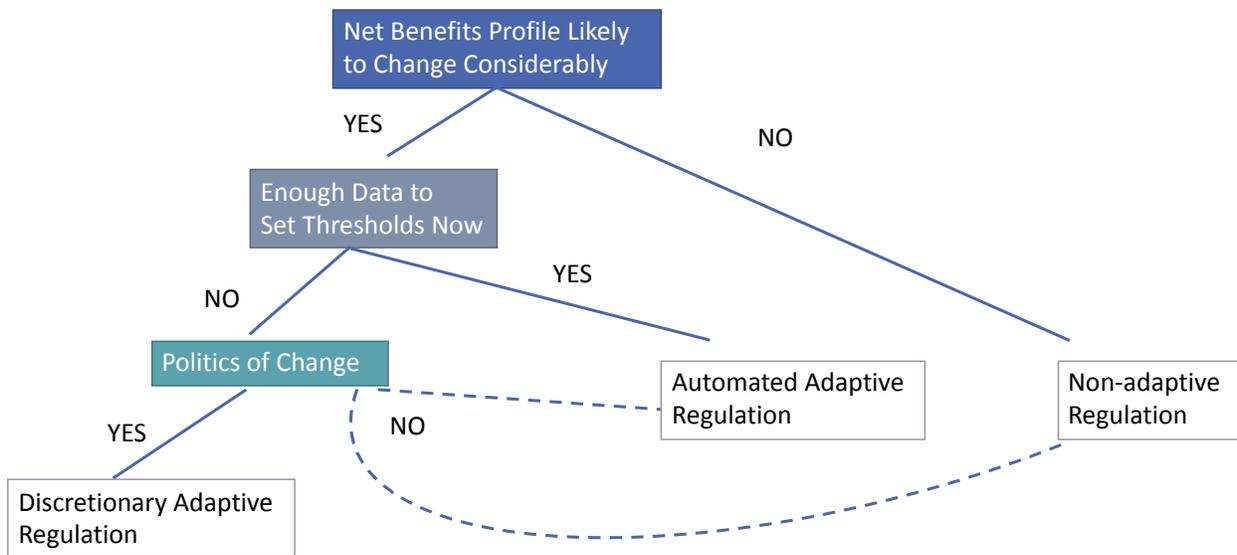


Figure 3: Decision Framework for Adaptive Regulation

4. Conclusions

Static regulation poses growing mismatches and social losses, especially in the face of rapid change in science, technology, or society. Adaptive regulation promises fewer such losses, but raises its own set of costs. Thus the net benefits of shifting from static to adaptive may vary by application, and may vary among the several different institutional mechanisms for implementing adaptive regulation. A framework for adaptive regulatory instrument choice is needed.

In this paper we have identified several adaptive instrument options, including:

- Policy variation – laboratory of decentralization, subsidiarity – learning from variation across jurisdictions
- Policy experimentation – laboratory of deployments with monitoring – learning from simulations, actual deployments, randomized trials

- Crisis response (unplanned adaptive) via ad hoc commissions of inquiry, or standing safety investigation boards
- Retrospective review, ex post Impact Assessment (IA) (ad hoc, vs. prospective commitment to do so)
- Sunset clauses (default of termination, after X years or threshold outcome reached, with possible renewal – on what criteria/showing) (or, default of continuation unless decision to revise/rescind)
- Repeated periodic policy reviews (e.g. every X years)
- Ongoing continual adaptive IA (with Adaptive RIA)
- Market-based incentives with banking/borrowing, price ceilings/floors
- Automated policy adjustment triggered by monitoring indicators
- Discretionary policy adjustment deliberated by an oversight body
- Adaptive licensing to targeted subpopulations, monitoring, updating/adding (e.g. for pharmaceuticals; automated vehicles; etc.)

And we have identified criteria for selection, including:

- Benefits
- Costs
- Ancillary impacts – co-benefits, countervailing risks; scope, breadth.
- Distributional equity (impacts on subgroups).
- Pace of change – scientific, technological, societal.
- Expected learning; benefits of reducing policy errors.
- Decision and information costs.
- Instability; costs of repeated or unanticipated policy changes.

- Institutional authority; power.
- Political economy; influence of anticipated adaptability on prospects for ex ante policy adoption, compliance, and on ongoing policy conflicts

Instrument choice for adaptive regulation needs to assess these instrument options and criteria in specific application areas. The choice among adaptive instruments confronts several likely tradeoffs, depending on key elements such as the frequency of review (posing the tradeoff of more rapid updating vs. greater instability), the scope of impacts to be assessed (posing the tradeoff of full social impact assessment including unintended ancillary impacts, vs. the costs of greater data collection and analysis), the decision maker (posing the tradeoff of automatic vs. discretionary – rapid and reliable, but based on a limited scope of impacts, vs. slower and more deliberate, but assessing a broader scope of impacts), and the implications for the political economy of regulation (such as easing initial adoption vs. undermining subsequent compliance).

References

- Adler, Jonathan H. 2016. Dynamic Environmentalism and Adaptive Management: Legal Obstacles and Opportunities. PERC Reports, reprinted from *Journal of Law, Economics & Policy* (Summer 2015), 11: 133-161.
- Aldy, Joseph E. 2014. "Learning from Experience: An Assessment of the Retrospective Reviews of Agency Rules and the Evidence for Improving the Design and Implementation of Regulatory Policy." Administrative Conference of the United States.
- Ayres, Ian, and John Braithwaite. 1994. *Responsive Regulation*. Transcending the Deregulation Debate. Oxford University Press on Demand.
[http://books.google.com/books?id=83XTEGB2mx4C&printsec=frontcover&dq=YRES+and+JOHN+BRAITHWAITE+RESPONSIVE+REGULATION+TRANSCENDING+THE+DEREGULATION+DEBATE+4+\(\(1992\)-\(Highlighted+jun+12+2015\)\)&hl=&cd=9&source=gbs_api](http://books.google.com/books?id=83XTEGB2mx4C&printsec=frontcover&dq=YRES+and+JOHN+BRAITHWAITE+RESPONSIVE+REGULATION+TRANSCENDING+THE+DEREGULATION+DEBATE+4+((1992)-(Highlighted+jun+12+2015))&hl=&cd=9&source=gbs_api).
- Balleisen, Edward J, Lori S. Benneer, David Cheang, Jonathon Free, Megan Hayes, Emily Pechar, and A. Catherine Preston. 2017a. "Institutional Mechanism for Investigating the Regulatory Implications of a Major Crisis: The Commission of Inquiry and the Safety Board." Chapter 17 in *Policy Shock: Recalibrating Risk and Regulation after Oil Spills, Nuclear Accidents, and Financial Crises*, edited by Edward J Balleisen, Lori S. Benneer, Kimberly D. Krawiec, and Wiener, Jonathan B. Cambridge, U.K.: Cambridge University Press.
- Balleisen, Edward J., Lori S. Benneer, Kimberly D. Krawiec, and Jonathan B. Wiener, eds. 2017b. *Policy Shock: Recalibrating Risk and Regulation after Oil Spills, Nuclear Accidents, and Financial Crises*. Cambridge, U.K.: Cambridge University Press.
- Benneer, Lori S., and Cary Coglianese. 2013. "Flexible Approaches to Environmental Regulation." In *The Oxford Handbook of U.S. Environmental Policy*, edited by Sheldon Kamieniecki and Michael E. Kraft. Oxford, U.K.: Oxford University Press.
- Birkland, Thomas A. 2006. *Lessons of Disaster: Policy Change after Catastrophic Events*. Washington, DC: Georgetown University Press.
- Botkin, Daniel B. 1990. *Discordant Harmonies: A New Ecology for the 21st Century*.
- Botkin, Daniel B. 1996. "Adjusting Law to Nature's Discordant Harmonies." *Duke Environmental Law & Policy Forum* 7: 25-38.
- Brunner, Ronald D. 2010. Adaptive Governance as a Reform Strategy, *Policy Sci.* 43: 301- .
- Burtraw, Dallas, Charles Holt, Karen L. Palmer, Anthony Paul, and William Shobe. 2017. "Expanding the Toolkit: The Potential Role for an Emissions Containment Reserve in RGGI." *Resources for the Future*.
- Carlson, Ann, and Dallas Burtraw, eds. Forthcoming 2018. *The Future of U.S. Energy Policy: Lessons from the Clean Air Act*. Cambridge, U.K.: Cambridge University Press.
- Carpenter, Daniel, and Gisela Sin. 2007. "Policy Tragedy and the Emergence of Regulation: The Food, Drug, and Cosmetic Act of 1938." *Studies in American Political Development* 21 (02): 149–80.
- Coffee, John. 2012. "The Political Economy Of Dodd-Frank: Why Financial Reform Tends To Be Frustrated And Systemic Risk Perpetuated." *Cornell Law Review* 97: 1019–82.
- Coglianese, Cary. 2013. "Moving Forward with Regulatory Lookback." *Yale J. on Reg. Online* 30: 57–57.

- Coglianesi, Cary & David Lehr. 2017. Regulating by Robot: Administrative Decision Making in the Machine-Learning Era, *Georgetown L.J.* 105: 1147- .
- Cosens, B. A., L. Gunderson, and B. C. Chaffin. 2018. "Introduction to the Special Feature -- Practicing Panarchy: Assessing legal flexibility, ecological resilience, and adaptive governance in regional water systems experiencing rapid environmental change." *Ecology and Society* 23(1): 4.
- Craig, R. K., A. S. Garmestani, C. R. Allen, C. A. Arnold, H. Birgé, D. A. DeCaro, A. K. Fremier, H. Gosnell, and E. Schlager. 2017. "Balancing stability and flexibility in adaptive governance: an analysis of tools available in U.S. environmental law." *Ecology and Society* 22(2): 3.
- Craig, Robin and J.B. Ruhl. 2014. "Designing Administrative Law for Adaptive Management." *Vanderbilt Law Review* 67: 1-87.
- Cropper, Maureen, Art Fraas and Richard Morgenstern. 2017. "Looking backward to move regulations forward: Rigorous ex post analyses can improve regulatory policies." *Science* 355: 1375-1376 (31 March).
- DeCaro, D. A., B. C. Chaffin, E. Schlager, A. S. Garmestani, and J. B. Ruhl. 2017. "Legal and institutional foundations of adaptive environmental governance." *Ecology and Society* 22(1): 32.
- Eichler, Hans-Georg, Kenneth Oye, et al. 2012. Adaptive licensing: taking the next step in the evolution of drug approval. *Clin. Pharmacol. Ther.* 91: 426–437.
- Farber, Daniel A. 1994. Environmental Protection as a Learning Experience, *Loyola L.A. L. Rev.* 27: 791- .
- Freeman, Jody & Daniel Farber, 2005. Modular Environmental Regulation, *Duke Law Journal* 54: 795-912.
- Graham, John D., and Jonathan B. Wiener, eds. 1995. *Risk vs. Risk: Tradeoffs in Protecting Health and the Environment*. Harvard University Press.
- Greenstone, Michael. 2009. Toward a Culture of Persistent Regulatory Experimentation and Evaluation, in *New Perspectives On Regulation* (David Moss & John Cisternino eds.).
- Gubler, Zachary J. 2014. Experimental Rules, *B.C. L. Rev.* 55: 129- .
- Holling, Crawford S. 1978. *Adaptive Environmental Assessment and Management*. John Wiley & Sons.
- IRGC 2017. International Risk Governance Council, Transatlantic Patterns of Risk Regulation: Implications for Trade and Regulatory Cooperation.
- IRGC 2018. International Risk Governance Council, The Governance of Decision-Making Algorithms (report of a workshop).
- Kahn, Matthew E. 2007. "Environmental Disasters as Risk Regulation Catalysts? The Role of Bhopal, Chernobyl, Exxon Valdez, Love Canal, and Three Mile Island in Shaping U.S. Environmental Law." *Journal of Risk and Uncertainty* 35: 17–43.
- Kaplow, Louis. 1986. "An Economic Analysis of Legal Transitions." *Harvard Law Review* 99: 509-617.
- Kaplow, Louis. 1992. "Government Relief for Risk Associated with Government Action." *Scandinavian Journal of Economics* 94: 525-541.
- Kaplow, Louis. 2003. "Transition Policy: A Conceptual Framework." *Journal of Contemporary Legal Issues*.

- King, Andrew A, and Michael J Lenox. 2000. "Industry Self-Regulation without Sanctions: The Chemical Industry's Responsible Care Program." *The Academy of Management Journal* 43 (4):698–716.
- Kroll, Joshua A., Joanna Huey, Solon Barocas, Edward W. Felten, Joel R. Reidenberg, David G. Robinson, and Harlan Yu. 2017. Accountable Algorithms. *U. Penn. L. Rev.* 165: -- .
- Lazarus, Richard J. 2009. Super Wicked Problems and Climate Change: Restraining the Present to Liberate the Future. *Cornell L. Rev.* 94: 1153-1233.
- Lindblom, Charles E. 1959. The Science of "Muddling Through," *Public Admin. Rev.* 19: 79- .
- Listokin, Yair. 2008. Learning Through Policy Variation, *Yale L.J.* 118: 480- .
- Lo, Andrew W. 2017. Adaptive Markets: Financial Evolution at the Speed of Thought.
- Maniloff, Peter, and Brian C Murray. 2009. "Allowance Price Containment Options for Cap-and-Trade Legislation." Nicholas Institute for Environmental Policy Solutions Working Paper 09-07.
- Marchant, Gary E. et al., eds. 2011. *The Growing Gap between Emerging Technologies and Legal-Ethical Oversight*. The International Library of Ethics, Law and Technology, Springer Science+Business Media B.V.
- Marchant, Gary E., Kenneth W. Abbott and Braden Allenby, 2013. *Innovative Governance Models for Emerging Technologies*. Edward Elgar.
- Maxwell, John W., Thomas P. Lyon, and Steven C. Hackett. 1998. "Self-Regulation and Social Welfare: The Political Economy of Corporate Environmentalism." *SSRN Electronic Journal* 6. <https://doi.org/10.2139/ssrn.135628>.
- McCray, Lawrence E, Kenneth A Oye, and Arthur C Petersen. 2010. "Planned Adaptation in Risk Regulation" *Technological Forecasting & Social Change* 77: 951–959.
- Morgenstern, Richard D. 2015. "The RFF Regulatory Performance Initiative: What Have We Learned?"
- Murray, Brian C, Richard G Newell, and William A Pizer. 2009. "Balancing Cost and Emissions Certainty: An Allowance Reserve for Cap-and-Trade." *Review of Environmental Economics and Policy* 3 (1):84–103.
- Percival, Robert V. 1998. "Environmental Legislation and the Problem of Collective Action." *Duke Environmental Law and Policy Forum* 9: 9–28.
- Pidot, Justin R. 2015. Governance and Uncertainty, *Cardozo L. Rev.* 37: 112- .
- Pizer, William A. & Brian C. Prest. 2018. Prices versus Quantities with Policy Updating. Draft of July 9.
- Pound, Roscoe. 1923. *Interpretations of Legal History* I.
- Repetto, Robert, ed. 2006. *Punctuated Equilibrium and the Dynamics of US Environmental Policy*. New Haven, CT: Yale University Press.
- Ribeiro, Daniel L. 2018. *Adaptive Regulatory Impact Assessment: Beyond the Foresight-Hindsight Divide*. SJD Dissertation, Duke Law School.
- Sabel, Charles F. & William H. Simon, 2011. Minimalism and Experimentalism in the Administrative State, *Georgetown L.J.* 100: 53- .
- Schmalensee, Richard and Robert N. Stavins, 2013. The SO2 Allowance Trading System: The Ironic History of a Grand Policy Experiment, *J. Econ. Perspectives* 27: 103- .
- Stavins, Robert N. 1998. What Can We Learn from the Grand Policy Experiment? Lessons from SO2 Allowance Trading, *J. Econ. Perspectives* 12: 69- .
- Sunstein, Cass R., 2014. The Regulatory Lookback, *B.U. L. Rev.* 94: 579- .
- Sunstein, Cass R. 2019. Sludge and Ordeals. *Duke L.J.*, forthcoming.

- Stewart, Richard B. 2014. Remediating Disregard in Global Regulatory Governance: Accountability, Participation, and Responsiveness, *American Journal of International Law* 108: 211- .
- Tarui, Nori & Stephen Polasky, 2005. Environmental Regulation with Technology Adoption, Learning and Strategic Behavior, *J. Env'tl. Econ. & Mgmt.* 50: 447- .
- Vogel, David. 2017. Policy Learning, in *Achieving Regulatory Excellence* (Cary Coglianese ed.)
- Wagner, Wendy, William West, Thomas McGarity & Lisa Peters, 2017. Dynamic Rulemaking, *N.Y.U. L. Rev.* 92: 182- .
- Weitzman, Martin L. 1974. Prices vs. Quantities. *Review of Economic Studies*, 41: 477-491.
- Wiener, Jonathan B. 1995. "Law and the New Ecology: Evolution, Categories and Consequences." *Ecology Law Quarterly* 22: 325-357.
- Wiener, Jonathan B. 1996. "Beyond the Balance of Nature." *Duke Environmental Law & Policy Forum* 7: 1-24.
- Wiener, Jonathan B. 1998. Managing the Iatrogenic Risks of Risk Management. *Risk: Health Safety and Environment* 9: 39-82.
- Wiener, Jonathan B. 2002. "Precaution in a Multirisk World." Chapter 32 in Dennis Paustenbach, ed., *Human and Ecological Risk Assessment: Theory and Practice* 1509-1531. New York: John Wiley & Sons.
- Wiener, Jonathan B. 2004. "The Regulation of Technology, and the Technology of Regulation." *Technology in Society* 26: 483-500.
- Wiener, Jonathan B. 2016. "The Tragedy of the Uncommons: On the Politics of Apocalypse." *Global Policy* 7: 67-80 (issue S1, May).
- Wiener, Jonathan B. 2017. "Risk Regulation and Future Learning." *European Journal of Risk Regulation* 8: 4-9.
- Wiener, Jonathan B. 2018. "Precautionary Principle," chapter in the *Encyclopedia of Environmental Law* (Michael Faure, ed.), volume on *Principles of Environmental Law* (Ludwig Kramer and Emanuela Orlando, eds.). IUCN and Edward Elgar.
- Wiener, Jonathan B. and Daniel L. Ribeiro. 2016. "Environmental Regulation Going Retro: Learning Foresight from Hindsight." *Journal of Land Use & Environmental Law* 32: 1-72.
- Wiener, Jonathan B. and Barak Richman. 2010. "Mechanism Choice," in *Research Handbook on Public Choice and Public Law*, Daniel A. Farber and Anne Joseph O'Connell, eds. Edward Elgar.
- Williams, Byron K. 2011. "Passive and Active Adaptive Management: Approaches and an Example." *Journal of Environmental Management* 92 (5):1371-78. <https://doi.org/10.1016/j.jenvman.2010.10.039>.
- Wuthnow, Robert. 2010. *Be Very Afraid: The Cultural Response to Terror, Pandemics, Environmental Devastation, Nuclear Annihilation, and Other Threats*. Oxford, UK: Oxford University Press.