

Chapter 13

A Conceptual Model of Planned Adaptation (PA)



Jesse Sowell

Abstract

- Technical and operational systems, especially large complex socio-technical systems, must continuously adapt to a changing context.
- Conventional rulemaking systems have, historically, been ponderous. Low clock-speed rulemaking systems are rarely equipped to keep pace when applied to govern higher clockspeed socio-technical systems.
- This chapter presents and characterizes four organizational complexes that have effectively implemented ‘planned adaptation’.
- Based on these cases, an initial general model of planned adaptation is presented: first, in terms of the first principles of where in the rulemaking system adaptation may be introduced; second, in terms of categories of events and triggers; and finally, in terms of the evaluative capabilities and capacities within and across organizations necessary to systematically plan to adapt.
- This general model (which can more generally be termed *systematic adaptation*) characterizes ad hoc adaptation, *planned* adaptation, as well as *ideal forms* of adaptation. These are used as comparators in what is referred to here as the ‘systematic adaptation space’.
- Discussion of four cases based on this model highlight instances of critical paths from ad hoc to *planned* adaptation, offering early insights that can lead to effectively designing organizational complexes that can effectively plan to adapt.
- The chapter concludes with a discussion of ongoing work—in particular, how this model will be used for theory-based sampling to identify cases of adaptation that can be used to further refine the model and its application to systematic adaptation by design.

J. Sowell (✉)

Department of International Affairs, The Bush School of Government & Public Service, Texas A&M University, College Station, TX, USA
e-mail: jsowell@tamu.edu

© The Author(s) 2019

V. A. W. J. Marchau et al. (eds.), *Decision Making under Deep Uncertainty*,
https://doi.org/10.1007/978-3-030-05252-2_13

289

13.1 Introduction

Historically, rulemaking has been a ponderous process. In contemporary societies, rulemaking must cope with combinations of uncertainties that not only arise when dealing with the natural environment, but also those uncertainties that arise from the built, technical environment. Rules—more accurately *rulemaking systems*—must be able to keep pace, incorporating credible knowledge assessments from relevant experts apace with the impact of technology on industry and society. In effect, rulemaking systems must be able to adapt. Rulemaking does not happen in a vacuum; rulemaking happens in organizations with vastly different resources, processes, capabilities, and capacities. The work described in this chapter builds on the existing literature on planned adaptation to offer a conceptual model, based on empirical cases, of the organizational apparatus necessary to execute adaptive policy recommendations.

The literature on decision analysis and its progeny—scenario planning, adaptive policymaking, Robust Decision Making, strategic options—have been the bread and butter of policy analysts since the 1950s. The notion of planned adaptation is a new and understudied aspect of rulemaking. The term was established by McCray et al. (2010), characterizing programs that “both revise rules when relevant new knowledge appears, and take steps to produce such improved knowledge”.

In contrast to work on decisionmaking in general and decisionmaking under deep uncertainty in particular, the object of planned adaptation is the character and context of the programs acting on new knowledge, *not* the underlying policy analysis. This is a subtle, but potentially important conceptual distinction, between the substance of the knowledge that comprises policy *recommendations* and the elements of the organization that must implement that recommendation. (This distinction is made clear in Chap. 14, which describes the implementation of Adaptive Delta Management in the Netherlands.)

In this chapter, McCray’s notion of planned adaptation is elaborated to explicitly characterize the apparatus necessary to systematically *plan for* adaptation. *To plan* focuses on resources and processes necessary to support the development of, recognize, and/or act on new knowledge:

1. What rules and mechanisms are in place to act on, elicit, consume, or refine policy-relevant information?
2. Who generates this new knowledge?
3. What is the organizational relationship of those generating new knowledge to those adapting the rules?

The process of adaptation comprises the set of resources, both economic and political, necessary to engender the recommended change. Taken together, planned adaptation speaks not only to the act of revising, but elaborates the permutations of organizational configurations necessary to plan for and implement stable, sustainable, systematic adaptation.

Studies of decisionmaking under uncertainty provide a rich analytic framework for evaluating information gaps, but rarely speak to the organizational and political

context in which triggers for evaluative action occur. To understand this context, concepts from the theory of rulemaking (Hart 1994) are used as a common conceptual vernacular for understanding adaptation in terms of the rules supporting that process. Hart's categories of rules facilitate disentangling substantive (primary) rules governing the *behavior* of a system from the categories of rules (secondary) used to *identify, evaluate, and change* rules when the rule system no longer effectively sustains system integrity. In general, primary rules represent the current understanding of how to govern the system—in effect a snapshot of a static policy analysis. Secondary rules describe (1) the evaluative capabilities necessary to transform new information about the system into prescriptive knowledge about the (mis)alignment of primary rules with system function and (2) the rules for identifying, creating, and maintaining the capacity necessary to change both primary and secondary rules as necessary, based on prescriptive knowledge.

Moving from abstractions to empirical mechanisms, based on the cases presented in this chapter, an initial set of variables is established for comparing instances of planned adaptation in terms of how secondary capabilities and capacities are used to adapt primary rules. Variables identified in this model characterize the factors that affect the development, implementation, and application of secondary rules. Variables fall into two categories: those that characterize sources of new information and those that characterize the rough organizational structure(s) supporting evaluative capabilities and capacities. Sources of new information are characterized in terms of *triggers* that signal that adaptation may be warranted, and the character of *events* producing the new information. Evaluative variables characterize the *timing* of evaluation relative to triggers and events; the *loci* of evaluative capabilities and capacities in the organizational complex; and how *coupled* the rulemaking principals are to the evaluative agents.

This chapter argues that the four cases presented here offer sufficient diversity in triggers, event context, and organizational configurations to lay the conceptual foundations for an initial model of systematic adaptation. In terms of the case studies presented here, these variables provide the foundation for concisely articulating adaptive capabilities and capacities of empirical organizational complexes. The model also admits “ideal” adaptive organizational complexes, such as Weber's notion of a high-functioning bureaucracy. Taken together, this model represents the “systematic adaptation space”. Mapping these variables facilitates comparing instances with little to no adaptive capabilities, ad hoc adaptation, and planned adaptation. More pragmatically, the analysis argues that, with additional cases, this model could provide the foundation for intentionally designing adaptive organizations.

Section 13.2 summarizes cases on delta management, particulate matter, air transportation safety, and Internet number delegation to provide comparative instances of organizations that have not only planned to adapt, but for which there are empirical instances of successful implementation of the adaptation. In Sect. 13.3, concepts from adaptive policymaking—in particular the character of triggers that signal an adaptation should take place—are synthesized with concepts from the theory of rule-making to contextualize the factors affecting resource allocation necessary to plan for adaptation. In Sect. 13.4, the variables identified are used to analytically compare

these cases and to identify ideal forms, effectively elaborating the range of configurations and contexts that characterize the systematic adaptation space. This systematic adaptation space provides an initial set of bounds that explicitly distinguish organizations with the capabilities and capacity to systematically and intentionally plan for adaptation from those that adapt ad hoc. Given these bounds, the chapter concludes by refining the notion of a systematically adaptive organization, offers a nascent framework for designing such an organizational configuration, and presents a sketch of future studies on planned adaptation.

13.2 Planned Adaptation Cases

Cases of successful, planned, systematic adaptation are not as common as one would hope. Four cases are discussed as the foundations of a conceptual model of systematic adaptation. In contrast to McCray's work, the conceptual unit of analysis in these cases is the rulemaking system, comprising one or more organizations collaborating to plan for adaptation. This unit of analysis includes state-based *and* non-state organizational complexes. The three state-based cases are: (1) a comparison of particulate matter standards development in the European Union (EU) and USA (Sect. 13.2.1); (2) flood risk safety in the Netherlands; and (Sect. 13.2.2) (3) civil air transportation safety standards development in the US (Sect. 13.2.3). The non-state case describes the private transnational regime that manages Internet protocol (IP) address delegation (Sect. 13.2.4).

Across these cases, key variables in systemic adaptation are introduced:

1. The degree of coupling of adaptive processes with substantive rules,
2. The provenance of events requiring change,
3. The timing of triggers signaling the potential need for policy evaluation and adaptation,
4. The loci of monitoring and evaluative capabilities, and
5. Social, political, and organizational incentives for investing in monitoring and evaluative resources.

In the cases, these variables are introduced in context; the variables are defined more formally in the section "Generalizing Elements of Planned Adaptation" (Sect. 13.3).

13.2.1 *Particulate Matter Standards*

Particulate matter (PM) standards in the EU and the USA comprise a canonical instance of planned adaptation. In both cases, the characteristic element of planning is a *periodic trigger*. In short, regulations limiting the emissions of particulate matter were introduced, with stipulations that combinations of modeling and monitoring would be used to re-evaluate these regulations after a given period of time. At a high

level, there are two classes of particulate matter considered in these discussions: PM_{10} and $PM_{2.5}$. The difference is in the size of the particles and their effects on air quality in terms of health impacts. PM_{10} establishes thresholds around particles with a diameter up to 10 micrometers; $PM_{2.5}$ establishes thresholds at a diameter of 2.5 μm . As an instance of planned adaptation, the periodic trigger demands a re-evaluation of the substance of the standard (or, very simply, whether the threshold is PM_{10} or $PM_{2.5}$) based on monitoring of particulate matter levels and reconsideration of the regulation based on additional scientific research on the effects of particulate matter.

One of the key findings stressed by Petersen et al. (2006) is the difference in the interpretation of science and monitoring data available to both EU and US standards bodies. In the case of PM standards, the USA, based on common scientific data, was quicker to introduce PM standards and was more precautionary.¹ Although the EU and USA “both use highly transparent, participatory process to inform policy decisions”, Petersen et al. (2006) attribute part of the nine-year lag to the review process. It should be stressed that the lag is not intended to say one organizational and regulatory process is better or more efficient. Rather it is a point of comparison; this work will build on to understand the loci of organizational, knowledge creation, and knowledge assessment (evaluation) capabilities in the adaptation of rule systems. This distinction speaks to what will be developed here as the capability to create, consume, and apply knowledge to anticipating change, and the need for subsequent adaptation based on established criteria. In the case of PM standards, Petersen et al. (2006) highlight that not only is the lag due to the review process, but that some of the misanticipation was rooted in misunderstanding (1) the impact of knowledge assessments themselves, such as the costs of detailed modeling in contrast to the observable benefits from these assessments and the impact on overall processes, and (2) how strictly standards would be enforced by the courts. In both of these instances, the character and loci of how knowledge and standards are evaluated and consumed affected both the adaptive process and the application of the standards themselves. In later discussion, this distinction will be elaborated in terms of the interaction between secondary rules (those used to change rules) and the application of primary rules (the substance of a standard or a regulation).

In the EU standards development process, the loci of adaptive capabilities are rather distributed, starting in the European Commission (EC), with implementation of standards by EU states. Petersen et al. (2006) describe the Clean Air for Europe (CAFE) program as the program that drove the development of the 2005 Thematic Strategy on Air Pollution. CAFE provided policy guidance. Technical expertise was contributed through the Technical Analysis Group (TAG), comprising guidance from

¹Typically the roles are reversed. The US regulatory system is considered generally less precautionary than EU regulation. US regulation is considered to be normatively liberal: regulations are typically driven by harms observed (and typically litigated) in a given industry, leading to particular rules safeguarding against those harms. In contrast, EU state regulations are considered to be socially protective: harms are anticipated and the attendant behaviors regulated, with derogations introduced when this precautionary approach is shown to create market inefficiencies in particular cases.

the CAFE secretariat with “consultants carrying out the technical analyses under specific contracts”. As noted by Petersen et al., “[t]he organization of this process can be seen as being top-down”. This science–policy interface is framed as “top-down”; in terms of organizational complexes, this will be referred to as a top-down configuration. In contrast, US regulation was driven by lawsuits—a more “bottom-up” configuration, injecting new knowledge into the adaptive process by those experiencing the costs of inefficient policies. As this work moves through the remaining cases and elaborates triggers driven by exogenous events and operational experience, it will become evident that this interface may be driven by feedback loops in top-down, bottom-up, or relatively lateral configurations within a given bureaucracy. Petersen et al. (2006) describe this interface in terms of “boundary work”. The distinction between actors engaging in risk assessment (attributed to science) and risk management (attributed to policy) is seemingly confounded when the distinction between stakeholders and experts is blurred. This distinction is rather fluid in the cases presented here—in particular the cases on regional Internet registries (RIRs), where the participants in an adaptive consensus process are stakeholders *and* operational experts with a knowledge base distinct from academic network engineers (science).

13.2.2 Delta Management in the Netherlands²

In the Netherlands, 55% of the country is at risk of flooding. Dutch present-day flood risk management is an instance of planned adaptation, whose roots stretch back to the coastal flooding disaster of 1953. During the multi-decade project that followed the flooding, substantive “knowledge about the necessary heights and strengths of dunes and dikes was generated ... [but] no system was put in place to systematically refine and update this knowledge and the resulting ‘hydraulic requirements’, engineering standards that correspond to the set safety standards for the different locations along the coast” (Petersen and Bloemen 2015). In terms of adaptation, the multi-decade response to the 1953 disaster was a systematic effort to improve flood safety, but, in terms of planning, it was an ad hoc response to an event.

Upon completion of the Delta Works project in 1989, planned adaptation was introduced to review safety standards and fund research necessary for generating new knowledge of the systems, risks, and investments in monitoring and data collection. The history of this project provides insight into the evolution of planned adaptation driven by periodic triggers. After the Zuiderzee flooding in 1916, the Rijkswaterstaat (the Directorate-General for Public Works and Water Management) evaluated whether flooding could happen again, and, based on the state of the art at the time, determined that it could not. In 1938, a civil engineer in the Rijkswaterstaat, Johan van Veen, challenged the assumption that no new knowledge would make a re-assessment worthwhile, initiating a study that ultimately led to the introduction of

²This case is described in detail in Chap. 14.

a plan with higher and newer dikes based on examination of local conditions. Unfortunately, this was too little too late, and “the country was ‘slapped in the face’ by the 1953 flood known in the Netherlands as ‘The Disaster’” (Petersen and Bloemen 2015).

Two factors contributed to the investment of political and organizational capital necessary to yield the Delta Works Plan. First, during World War II, some portions of the Netherlands were intentionally flooded as a military defense against the Germans. “After the war, the Dutch were able to quickly regain the land (sometimes after surmounting huge technical difficulties) and they started to believe more and more that many desired physical and social changes in the Netherlands could be effected by choosing the right government policies and engineering approach (*‘maakbaarheid’*)” (Petersen and Bloemen 2015).

Experience with rebuilding provided evidence that, with a concerted political and engineering effort, substantive progress could be made. The second factor was the fallout from “The Disaster”: “Though the number of casualties was not very high as compared with earlier disasters (1,836 people drowned as compared to tens of thousands in earlier floods in the 1500s for instance), the effects were traumatic—at the individual level, for the Netherlands as a country, and for the coastal engineering profession” (Petersen and Bloemen 2015).³

Taken together, these two historical events contributed social capital, and perhaps a new impetus for the coastal engineering profession to improve its engineering and technical capabilities to better serve the public good. In effect, it sets the stage for investing in the organizational, knowledge creation, and monitoring capabilities that ultimately produced the Delta Works with, among others, the storm surge barriers in the Oosterschelde and the Nieuwe Waterweg.

In December 1995, the *Wet op de waterkering* (Water Defense Act) introduced statutory requirements for flood safety levels and “a five-year review cycle of the hydraulic requirements associated with those *fixed* safety levels” (Petersen and Bloemen 2015; emphasis added here). This is a canonical instance of periodic planned adaptation. In the first iteration, the review period was five years. Note above the emphasis added to *fixed*—during the review period, standards levels remain fixed. Upon (periodic) review, these standards, in particular local standards, may be updated to reflect new knowledge of the system.

In terms of the model of systematic adaptation being developed here, the experiences driving adaptation are a combination of flood events in 1916 and 1953 and the experience with both rebuilding after World War II and ad hoc adaptation in the first Delta Plan. To be clear, experience is the aggregate knowledge produced over *a series of events*; experience should not be mistaken as an analytic construct for scoping events or triggers. The trigger is the periodic review. Adaptation comprises the series of local updates that, after review, apply new knowledge to safety standards and, subsequently, their implementation.

The Dutch government constituted a second Delta Commission in 2007. The reasons: outdated flood safety standards (reflecting changes in population and value

³Petersen and Bloemen (2015) attribute this to Bijker (2007).

of investments in the early 1960s), a backlog in maintenance of dikes, and growing concerns about the possible consequences of the changing climate and sea level rise. In line with the advice of the Commission, the Delta Programme was established. The Delta Programme combined existing insights into water management and in dealing with uncertainty into an approach named Adaptive Delta Management (see Chap. 14).

Adaptive Delta Management (ADM) is based in part on adaptation pathways (see Chaps. 4 and 9). An important aspect of working with adaptation pathways is that the moment of adaptation of existing plans is not fixed, but is the result of an analysis of monitoring and modeling results on both the effectiveness of the measures being implemented and to (actual and expected) changing external conditions. In addition, two pre-fixed adaptation moments are built into the Delta Programme: (1) the overarching Delta Decisions (policy frameworks) and regional strategies are systematically reviewed every six years; and (2) the flood protection standards are reviewed every 12 years. Research programs like the Dutch National Knowledge and Innovation Program “Water and Climate” (NKWK) have a major role in generating the needed information. So the locus of this knowledge generation is the combination of research and lessons from adaptation.

In terms of the social and political capital necessary to sustain planned adaptation, the Netherlands is unique in that it has a distinct social awareness of the importance of water defense: This has been a part of their culture for centuries. Contrast this to the accrual of political capital for other cases: particulate matter and civil aviation safety are comprehensible by the politicians, but are not nearly as prominent a fact of life to individuals as water defense is to the Dutch. In terms of Internet numbers, most users are at best vaguely aware of what an IP address is, much less the impacts of their delegation, which is many layers of abstraction away from their day-to-day navigation of the World Wide Web (WWW—the application layer). It is easy to categorize this form of flood safety management as a somewhat common instance of periodic planned adaptation for rare events, but the difference in the sources of social and political capital make it a valuable instance that hints at an ideal form, where one finds a harmonious confluence of public social, political, and organizational capital necessary to invest in planned adaptation.⁴

13.2.3 *Air Transportation Safety*

In the USA, the Federal Aviation Administration (FAA) is responsible for “regulat[ing] safety in commercial air transportation and other modes of civil transportation” (McCray et al. 2010). In addition to a “static” license certifying airworthiness,

⁴Harmonious is used in the sense of Keohane (1984). In the ideal form, incentives are already aligned. Clearly in the Dutch case, these aligned incentives exist, but planned adaptation serves as a coordination mechanism that fosters the development of capabilities, and the deployment of capacity necessary for periodic review and the subsequent updating of standards.

the FAA has invested in a body with (1) the technical and operational capabilities necessary to evaluate accidents, (2) the capacity to investigate incidents as they occur, and (3) an established channel to provide feedback to update safety regulations. In terms of the language established here, this body has the necessary deep knowledge of airplane mechanics and broader aviation system operations. In contrast to the particulate matter case, where the majority of the knowledge brought to bear is based on science, in this case, the majority of knowledge brought to bear is technical and operational. This body of knowledge builds on scientific knowledge, but the utility of this evaluative and investigative body is its knowledge of how these systems operate on the ground, *in situ*.

This investigative body, the National Transportation Safety Board (NTSB), is a case of planned adaptation that is, in contrast to the previous two cases, driven by exogenous, stochastic events (accidents), not a periodic review driven by information and knowledge generated by an explicit event monitoring and/or research programs rooted in conventional scientific research. The NTSB deploys “go-teams” that are quickly mobilized to examine every civil aviation accident shortly after those incidents occur. In the terms established here, this is an investment of organizational capital to sustain an investigative capacity. The capability comprises the knowledge necessary to investigate an incident, evaluate the causes in the context of existing safety regulations, and recommend necessary, yet tractable, changes to safety regulations. The NTSB is an investment in capabilities, complemented by the appropriate investment in the capacity necessary to leverage those capabilities.

The result is a relatively tight feedback loop that leverages this combination of capabilities and capacity. As per McCray et al. (2010): “‘planning’ ... is the provision, in advance, of ample investigatory capacity ... to enable diagnosis of problems when they arise”. The adaptation feedback loop is animated by a body that advises, but is independent of, the FAA. The NTSB’s capabilities are triggered by the observation of exogenous events. Investigations, facilitated by preexisting capacity, generate the information and knowledge necessary to recommend adaptations.

The history of the NTSB speaks to the distinction between stakeholders and experts. Early on (until 1974), the NTSB was housed within the Department of Transportation (DOT). After an incident in which its recommendations were resisted, Congress ended the NTSB’s dependence on the FAA, making it an independent body. “While FAA personnel do participate in some NTSB investigations, they are pointedly excluded from NTSB deliberations to determine the ‘probable cause’ of civil aviation accidents” (McCray et al. 2010). In effect, establishing the independence of the NTSB created a distinct line between experts recommending regulatory adaptation and the stakeholders that will be affected by that change.

Since then, the NTSB has extended its scope beyond “mere disaster response”. The NTSB maintains a database of non-military incidents, supplementing individual investigations with efforts to identify patterns that may not be evident in individual cases. Under this remit, it investigates “near misses” and provides a list of “most wanted” safety improvements, effectively focusing public attention on those areas where technical or policy innovation is needed. This is another instance where we see a distinction in capabilities (research capabilities) and an investment in the capacity

necessary to support those application of these capabilities (accident database and supporting human resources).

As an interface in this organizational complex, this is an instance of lateral knowledge sharing between nominally peer agencies. The DOT/FAA is the stakeholder consuming new knowledge and updating rules. The NTSB is the expert community. When Congress made the NTSB independent of the DOT/FAA, it invested political capital (based on public demand), effectively imbuing the NTSB with the ability to enhance its capabilities (expert knowledge, investigative processes, understanding of near misses) and capacity (resources supporting ‘go-teams’, accident database), as it sees fit, in service of its investigative and evaluative functions.

13.2.4 Internet Number Delegation

The previous three cases describe the organizational capabilities supporting adaptation in conventional government agencies. Internet number delegation is managed by a transnational collection of non-profit organizations, referred to as the Regional Internet Registry (RIR) system.⁵ The RIR system comprises five regional registries (RIRs).⁶ Each RIR is a non-profit, membership-based organization tasked with delegating Internet Protocol (IP) addresses used for routing traffic from one Internet host to another.

IPv4 addresses are a finite, scarce resource, provisioned as a contiguous series of 2^{32} integers. The community often refers to these simply as “numbers”. IP addresses are used to uniquely identify Internet hosts in the Internet’s routing system. An Internet host is any device capable of Internet communication, comprising desktop computers, laptops, tablets, mobile phones, elements of the infrastructure such as routers and switches, servers used to host cloud computing services. Early on, the Internet operations community recognized the need to develop policy to order the delegation of this scarce resource, adapting variants of the consensus process used by the Internet Engineering Task Force (IETF).⁷

⁵The majority of this section is based on Chap. 4 of Sowell (2015), a study on the governance of RIRs and the consensus-based knowledge assessment processes these organizations use to develop resource delegation policy.

⁶The five RIRs are:

- American Registry of Internet Numbers (ARIN), covering North America and parts of the Caribbean
- Latin America and Caribbean Network Information Center (LACNIC), covering Latin America and parts of the Caribbean
- Réseaux IP Européens (RIPE) Network Coordination Centre (NCC), covering Europe, the Middle East, and Russia
- Africa Network Information Center (AFRINIC), covering Africa
- Asia Pacific Network Information Center (APNIC), covering Asia Pacific

⁷See Resnick (2014) for the modern articulation of the IETF consensus process.

In terms of policy development, (most) of the stakeholders in the RIR community are also numbers and routing experts. This community comprises network operators⁸ with deep expertise in the day-to-day operations of the technical systems and protocols that facilitate Internet communication. Network operators are an instance of what Haas (1992) refers to as an epistemic community: “An epistemic community is a network of professionals with recognized expertise and competence in a particular domain and an authoritative claim to policy-relevant knowledge within that domain or issue area.”

Conventional epistemic communities comprise professions such as lawyers, physicians, and academics, all of whom derive a substantive portion of their initial training, credibility, and foundational knowledge base from formal schooling. Operator communities differ; the training, credibility, and knowledge base is largely derived from operational experience managing a production Internet infrastructure. The immediate stakeholders—those whose firms depend on number delegations for their value proposition—form a superset of the experts actively evaluating and adapting number policy. Although there are certainly varying degrees of expertise in the larger numbers community, the collective body of numbers stakeholders has the necessary depth of knowledge regarding the role of these resources and how they are used in the operational environment to evaluate and adapt number delegation policy. As such, while those that animate the policymaking process are typically at the upper end of this spectrum, the stakeholder collective is itself an epistemic community.

Consensus-based knowledge assessment is at the heart of number resource policy development processes. Here, consensus-based decision making is used to evaluate and adapt (common) resource management rules to cope with uncertainties stemming from changing resource demands and patterns of use in the broader Internet industry. Resource policy is created by the community as a response to changes in operational use that warrants adaptation to align policy to operational realities. A resource policy proposal is crafted by one or more community members, drawing on their operational expertise, but acting in their role as stakeholders. That resource policy proposal is then evaluated by community experts and RIR staff for fairness; technical, operational, and economic feasibility; and operational implications for the RIR. This evaluation takes place as part of the consensus process. The four phases of the RIR’s consensus process are described below.⁹

Phase 1: Problem Identification

When members of the numbers community identify a policy inefficiency, the community must determine if it is legitimately within the scope of RIR number resource

⁸Here, “network operator” refers to an individual engineer working for a particular firm, not the firm itself.

⁹The following four phases describe the core tenets of consensus-based decisionmaking across the five RIRs. It should be noted that the generalization presented here does not elaborate the process-specific nuances in each RIR.

policy.¹⁰ Distinguished members of the policy development process (PDP),¹¹ referred to colloquially as “policy shepherds,” evaluate the problem statement to determine legitimacy. Problem identification involves an evaluation of the event—whether the implications warrant adaptation. Here, an event is a change in patterns of number use that warrants a change in delegation policy. Such events driving change are exogenous (typically driven by broader changes and innovation in Internet infrastructure) that are, in turn, a response to the demands from emerging patterns of use at the application layer.¹² The stakeholder that brings this event and its (perceived) implications to the attention of the broader community¹³ is the trigger for event evaluation and policy adaptation. In this sense, the event is exogenous, experienced by stakeholders that make up the numbers community, one or more of whom trigger a policy change.

Phase 2: Active Consensus

Where problem identification identifies whether a problem is legitimate and in scope, active consensus evaluates the proposal in terms of demand among operators experiencing inefficiencies, the quality of the proposal and the appropriateness of the proposed solution. Active consensus is the primary locus of policy adaptation. Within a given policy proposal, “adaptation step 0” is the initial proposal itself. Subsequent adaptations (steps 1 . . . n) occur as members of the numbers community make utterances that suggest credible modifications to the proposed solution. In that sense, active consensus is a process of knowledge assessment evaluating adaptations within the context of a given proposal.

Within the numbers community, politically fungible voting is eschewed as a knowledge assessment process.¹⁴ Under the family of consensus processes, asserting (voting) “No, I do not agree,” is insufficient. Credible contributions require a rationale that, through evaluative policy dialogues that encourage constructive conflict, either fit or update the operational epistemic communities’ authoritative image of number resource dynamics. Shallow agreement via a marginal “51–49 victory” is not consensus. In this transnational resource management space, there is no hegemon to enforce the rules. Rather, mutual enforcement requires credible commitments from participants that, as both architects of the rules and subjects of those rules, must both reason about how to adapt, adhere to, and enforce resource policy. As such, consensus often demands that >70% of credible contributors agree, building com-

¹⁰In ARIN, the Advisory Council (AC) evaluates initial policy proposals for scope and technical merit. In the RIPE and APNIC regions, the specific working group (WG) chairs evaluate policies based on the WG mandate.

¹¹AC members in ARIN; policy forum chairs in other regions.

¹²The canonical OSI model has seven layers, but pragmatically speaking only the infrastructure layers (1–4) and the application layer (7) are used in operational engineering environments.

¹³To be more explicit, this policy entrepreneur is likely one of a number of community members that has recognized the problem, but, as a policy entrepreneur, is incentivized to act on this by engaging the policy development process.

¹⁴Recall that the RIR’s consensus process is derived from the IETF consensus process. Voting is eschewed in that epistemic community, as well.

mitment by reconciling technical and operational critiques with dissenting minority contributors.

Iterative adaptation continues until all contestations have been addressed. Contestations may be integrated or dismissed. A contestation is integrated if the modification of the solution is itself not contested. If the modification suggested by a (credible) contestation is contested, the process of exploring the solution space continues until an uncontested solution is identified. When all contestations have been resolved, active consensus has been achieved and the process moves on to passive consensus.

Phase 3: Passive Consensus

Consensus does not require participation by every member of the numbers community. Passive consensus provides the opportunity for those that may not have had the time to follow each incremental change to weigh in on the product of active consensus. Passive consensus serves as a consistency check—a penultimate evaluation of the proposal that ensures uncertainties and/or inconsistencies that may have crept in during an iterative active consensus are not overlooked.

The character of passive consensus differs from active consensus. Active consensus requires both “active” support of a solution and the resolution of contestations. In contrast, silence on the proposal (the absence of contestation after a designated review period) is sufficient for achieving passive consensus. Over the course of the active consensus process, a policy proposal may have experienced a number of major and/or minor revisions. A minor contestation in the passive consensus process, such as a clarification of language, may be dealt with immediately in the passive consensus process. “Easy” cases, such as language cleanup or minor modifications, rarely see substantive change during the passive consensus process. In contrast, “hard” proposals are typically instances of larger policy issues, such as routing security, transfers, or anti-abuse that defy modularization and may have inconsistencies even after substantive efforts at active consensus. If a substantive contestation arises during passive consensus, the proposal may return to the active consensus phase for additional evaluation and revision.

Phase 4: Evaluation of the Consensus Process

A review of the consensus process itself is the last step of the PDP. Active and passive consensus is “called” by the shepherds of that particular proposal, indicating that phase is complete. This is a point of discretion by the consensus facilitator. The review phase is a check on that discretion.

Evaluating the substance of a proposal *is not* the objective of the review phase.¹⁵ The objective of the review is to ensure:

¹⁵In all but the RIPE region, the membership elected board of the RIR reviews policy proposals that have reached consensus. In the RIPE region, although issue-specific WG chairs determine consensus for policies developed in their WG, the collective of WG chairs evaluates whether the PDP was followed.

1. the RIR's PDP was followed, and
2. the policy does not create undue legal risk or obligations for the RIR.

For instance, in the ARIN region, the Board of Trustees of ARIN evaluates the draft policy in terms of “fiduciary risk, liability risk, conformity to law, development in accordance with the ARIN PDP, and adherence to the ARIN Articles of Incorporation or Bylaws,” (ARIN 2009). Review is intended to be final check on the consensus decisionmaking process and feasibility of the policy before moving to implementation. In terms of credible knowledge assessment, with the exception of instances where the board identifies a substantive risk to the RIR itself, the board does not subvert the credible knowledge assessment of the community.

Loci of Number Delegation Capabilities

RIR's and numbers community's investment in sustaining the consensus process is an investment in adaptive capabilities and supporting organizational capacity. There are two loci:

1. The members of the number community that evaluates of the substance of numbers policy constitute adaptive capability.
2. RIR staff supporting the infrastructure facilitating this process, performing impact analyses for RIR implementation, provides organizational capacity.

In terms of the organizational complex, this is a bottom-up interface in which an epistemic community is the locus of evaluative capabilities and the RIR (as an organization) provides organizational capacity (including coordination resources and process support) in which these evaluations can play out. Systematic adaptation through the consensus process sustains and extends the knowledge base of the operational epistemic community, in part by codifying that knowledge into resource policy.

13.3 Generalizing Elements of Planned Adaptation

As noted in the introduction, McCray, Oye, and Petersen's notion of planned adaptation is fundamentally about updating rule systems. There is subtle difference between (1) the knowledge necessary to make substantive policy decisions, and (2) the rulemaking apparatus, capabilities, and authority, necessary to change potentially entrenched rules. Starting with Hart's notion of primary and secondary rules, this conceptual model generalizes and explains the factors at play when adapting the rule systems presented here. Rules are ultimately evaluated in terms of their fit—how new knowledge of the system can be used to evaluate whether a set of rules continues to satisfy obligations to system integrity.

In the abstract, such a model is at best an elegant thought experiment. To make this model pragmatic, the notions of triggers, events, and loci of evaluation in terms of monitoring and knowledge assessment are formalized. Triggers are a known class

of mechanisms in adaptive rulemaking. Here, the notion of a trigger is the mechanism linking *events* identified as potentially warranting adaptation to the adaptive process itself: an evaluation of the substantive rules governing system behavior and the potential rules that provide guidance for how the adaptation should proceed. Categories of triggers serve as the bridge between an abstract model of adaptation and observed instances of adaptation discussed in the previous section. Events provide information that, upon assessment, offer an opportunity to adapt. Assessing an event demonstrates an evaluative capability to generate new knowledge of the system from information gleaned from an event, but not necessarily the capacity and capabilities necessary to *adapt the attendant rules*. Note the distinction between the capability to evaluate the substance of the event and the capacity and capabilities necessary to change the rules. The question for a model that aspires to be both analytic and evaluative then becomes: what are the necessary organizational and political capabilities necessary to consistently and reliably act on triggers and the substance of policy analyses?

Cases selected for this chapter are a comparatively interesting mix of decision-making processes, organizational structures, and incentive structures that provide a first view into the efficacy of the conceptual vernacular established in this section. That said, this is a relatively small convenience sample. The conceptual model presented here is a canonical case of contingent theory building.¹⁶ While limited, the concepts established here provide the basic building blocks for a first articulation of the systematic adaptation space—the range of configurations and contexts that sustainably facilitate systematic adaptation.

13.3.1 *Disentangling Primary and Secondary Rules*

Hart (1994) describes rule systems as the union of primary and secondary rules. *Primary rules* are those “thought important because they are believed to be necessary to the maintenance of social life or some highly prized feature of it”. Primary rules describe what behaviors are appropriate and what behaviors are prohibited. For instance, a prohibition on stealing is a primary rule. It is a rule that maintains social order, it ensures that the system (in this case a society) functions as one expects and that property rights are respected.

Much of the work on adaptive policymaking is about establishing primary rules that describe how to sustain system function and integrity. In the cases presented here, primary rules describe the appropriate (safe) levels of particulate matter, what is necessary to sustain appropriate flood protections, the body of flight safety rules, and the rules that articulate number delegation rights and how those rights are conferred. These are the rules that determine either how a system should be monitored and sustained, or how the system itself is to function. The approaches and tools of policy analysis described in the previous parts of this book are one conventional means by

¹⁶Contingent theory building is used in the sense of George and Bennett (2005).

which an analyst may systematically evaluate and develop the knowledge informing primary rules for maintaining a given system.

In contrast, *secondary rules* are about how to recognize, create, maintain, and adjudicate primary rules. Secondary rules are “all about [primary rules]; in the sense that while primary rules are concerned with the actions that individuals must or must not do, these secondary rules are all concerned with the primary rules themselves. They specify the ways in which the primary rules may be conclusively ascertained, introduced, eliminated, varied, and the fact of their violation conclusively determined” (Hart 1994).

Adaptive processes must strike a balance between the canonical objective of regulation (creating stability) and the potential for chaos if feedback loops trigger changes too frequently. Conceptually, primary and secondary rules give us a starting point for decomposing adaptive rulemaking processes. Consider the function of secondary rules above in terms of how primary rules may be “ascertained, introduced, eliminated, varied, and the fact of their violation conclusively determined” (Hart 1994). In terms of adaptation, ascertaining a rule’s fitness can be framed as part of the path selection process: given a particular trigger, what rule is most appropriate to cope with the conditions at hand? If no rule is fit, the introduction and elimination of rules represent the action that is taken to update a particular rule system—the act of actually adapting to introduce rules more appropriate, or better fit, to the better understanding of a system’s context, as highlighted by the evaluation of an observed event.

Hart also provides a notion of obligation. In this work, the obligation is to the efficiency and efficacy of the system at hand. In that sense, one could say that a rule will not be recognized (it will be determined to have poor fitness for the context) when it no longer fulfills its obligation to sustaining system integrity.

The secondary rules function and role provide the foundation for analyzing how decoupled primary and secondary rules are in a particular adaptive rulemaking system, and the implications of that decoupling for the performance of the system. In the next subsections, Hart’s three types of secondary rules (recognition, change, and adjudication) are presented in terms of how they help explain systematic adaptation.

Rules of Recognition

Hart (1994) argues that “the simplest form of remedy for the uncertainty of the regime of primary rules is the introduction of what we shall call a ‘rule of recognition’... what is crucial is the acknowledgment of reference to the writing or inscription as authoritative, i.e. as the proper way of disposing of doubts as to the existence of the rule”.

Simple rules of recognition may refer to an *authoritative* set of rules—in early societies an inscription on a monument, or in later systems, a particular constitutional document. In more sophisticated systems of rules, rules of recognition provide some “general characteristic possessed by the primary rules” such as “having been enacted by a specific body, or their long customary practice, or their relation to judicial decisions”.

Operationalizing recognition in terms of adaptation makes points of contention faced by both the creation of, and sustained application of, adaptation apparent. First, rules of recognition reduce uncertainty by highlighting precisely what characterizes a primary rule. In the context of an adaptive process, rules of recognition can also serve as a means to articulate scope—general characteristics can be framed in terms of the system being adapted. “Authoritative” speaks to the source of the rule. In this context, the “specific body” enacting the rule may well be a form of expert committee or epistemic community that has the necessary access to the system, knowledge, and capabilities to effectively develop rules that shape and/or sustain system behavior. Each of the cases in the previous section calls out those expert communities in terms of their evaluative capabilities.

Another operationalization of a characteristic that may be used in conjunction with the source of the rule is the *state of the art* of the rule. State-of-the-art speaks to how up-to-date the information informing a given rule is. Returning to McCray’s definition of planned adaptation, rules of recognition are an assertion that primary rules should be based on the best knowledge available on a given system. This seems like an obvious statement, but it serves an evaluative function that will serve as the foundation of a trigger for change. The criteria set by rules of recognition contribute to the broader objective of ascertaining whether a given rule should be in effect. A consequence of that evaluation is the choice to eliminate a given rule, introduce a new rule, or adapt the rule to reflect updated information.

It is important to stress that rules of recognition establish authoritative, not evaluative, criteria. Such criteria establish what it means for a rule to have an authoritative source—whether a rule is within the scope of the system. Recognition does not indicate *how* the rules should change, nor do they speak to the *substance* of primary rules. Speaking to the substance is more traditionally the domain of substantive policy analysis. The question then becomes, how do we systematically adapt the system when we recognize primary rules are no longer appropriate?

Rules of Change

Secondary rules of change “remedy . . . the static quality of the regime of primary rules . . . [by] empower[ing] an individual or body of persons to introduce new primary rules for the conduct of the life of the group, or of some class within it, and to eliminate old rules” (Hart 1994).

Rules of change provide the *mechanics* for adding new rules, updating existing rules, or removing rules that no longer satisfy obligations to system integrity. Depending on the rule system, those mechanics may be specialized to the regime in which the rule system is embedded or may be part of a more general rulemaking or legislative process. For instance, in the cases of particulate matter and flood risk, evaluative criteria are updated periodically, but the changes themselves are still part of existing regulatory processes. In contrast, in the case of airline safety, the NTSB is based on an established process of feeding information into the FAA to update safety standards. Moreover, the NTSB was explicitly given greater independence after allegations that the FAA and the White House may have undermined recommendations. Number policy in the RIR system illustrates an instance where rule change is almost

completely endogenous to the system: evaluation of primary rules is triggered by the constituency as they recognize failures of primary rules to the obligation to system integrity. Returning to McCray's initial definition, not only is the trigger endogenous, but the knowledge generation process is the factor driving change.

Rules of Adjudication

Hart frames the last set of secondary rules in terms of efficiency—in particular, how to move from primary rules enforced by social pressure, to vesting the authority for adjudicating primary rules into particular actors. Rules of adjudication comprise “secondary rules empowering individuals to make authoritative determinants of the question whether, on a particular occasion, a primary rule has been broken” (Hart 1994).

Like the earlier secondary rules, adjudication hinges on the notion of what constitutes authoritative (here “authoritative determinants”). Rules of recognition designate the authoritative bodies, as noted above, but rules of adjudication provide evaluative procedures.

Hart (1994) elaborates the linkages among secondary rules, as follows: “Besides ... resemblances to other secondary rules, rules of adjudication have intimate connections with them. Indeed, a system which has rules of adjudication is necessarily also committed to a rule of recognition of an elementary and imperfect sort. This is so because, if courts are empowered to make authoritative determination of the fact that a rule has been broken, these cannot avoid being taken as authoritative determinations of what the rules are”.

Here, Hart elaborates the link between rules of adjudication and recognition relative to courts that this work makes for the authoritative experts on the complex system to which policies must adapt. Conceptually, these are similar—an organization is established to systematize these processes. The pragmatic evaluative procedures are where these will differ, trading legalistic analysis for the delegation of decision-making authority to actors that hold epistemic authority in a particular domain of complex systems. In particular, adjudication differs from recognition of scope. Adjudication evaluates the performance of the rule, whether it is up-to-date, and, ultimately, whether it should be in effect under the current conditions.

Rules of adjudication speak to the criteria delineating *whether* one should change the rules or not, whether the change is warranted by system performance, or the implications for stakeholders (costs, efficiency, distributional effect, etc.). In the next subsection (Sect. 13.3.2), the entry point for recognizing and adjudicating the efficacy of a rule will be operationalized in terms of triggers and event evaluation. In effect, a trigger is a signal—a recognition of an event. That event provides new information about the performance of a rule system. For example, a periodic trigger is established by an authoritative (recognized) body to evaluate event monitoring information (particulate matter). The trigger is the entry point to an analysis under the adjudication function.

Taken together, secondary rules of change are one indicator of adaptive capability. They are durable procedures for systematically evaluating whether a policy

or set of policies, continue to meet obligations to system integrity. For each of the systems evaluated here, this organizational construct is the locus of permanent or semi-permanent authority to adjust rules to suit changes in the environment (like a judiciary in conventional governmental arrangements). Moreover, adjudication procedures will differ substantively. For instance, in conventional bureaucracies, domain experts contribute to the evaluation of the efficacy of rules (such as the EPA's periodic evaluation of particulate matter policy, or the event-driven evaluation of airline safety regulations). Expert adjudication is not limited to dedicated analysts, though. In the RIRs, the relevant experts are the engineers that enjoy the efficiencies, or suffer under the deficiencies, of number delegation policy; these actors serve as both monitor and change agent through RIRs' policy development processes. Each of these instances highlights not only the actors, but the triggers for policy change (the topic of the next subsection).

13.3.2 *Triggers and Events*

Adaptation here means processing system feedback—leveraging organizational capacity to transform information about the system into knowledge that informs the rules an organization uses to sustain system integrity and function. This subsection elaborates a conceptual model of triggers, events, and evaluative information processing and knowledge assessment that (ideally) culminates in effective adaptation of a set of rules. Definitions presented here bridge the substance of policy analysis and the adaptation of Hart's rules to policy adaptation. As noted earlier, rules of change describe how to change primary and secondary rules. They are the mechanics of change within an organizational process, but do not speak to the evaluative component. Triggers, events, and information processing map to the organizational complex's capabilities and capacity created by investing political capital in bodies with the expertise and authority necessary to effect rules of recognition and rules of adjudication. Rules of recognition authoritatively designate those that should monitor triggers and the structure of those bodies. Rules of adjudication authoritatively designate two elements of the evaluative process:

1. When a behavior deviates from the current definition of system integrity (i.e., when the current rules are no longer fulfilling their obligation to system integrity);
2. The process experts should use to evaluate the information inherent in an event to either select a more appropriate rule or, in some cases, create a new one.

In what follows, we elaborate these processes through systematic definitions of the types of triggers, events, and evaluation observed in the cases.

Triggers

A trigger *signals* when a change that potentially affects obligations to integrity has occurred; i.e., a trigger is the means of recognizing the potential need to adapt. (This is practically the same definition of trigger (or Adaptation Tipping Point) that is used in

the various DMDU approaches elsewhere in the book.) In the many uses of the term, a trigger is caused by an event, but that is not necessarily always the case. A trigger may also be the consequence of ongoing evaluations of events—this distinguishes reactive triggers from proactive triggers. Triggers activate an evaluative (adjudicative) process that informs whether rules of change should be invoked. *Proactive* triggers initiate an evaluation based on information that has already been collected by some monitoring process. Reactive triggers are more akin to the prevailing use of the term. *Reactive* triggers activate both an information gathering process and evaluative process in the face of a previously unobserved, unanticipated, or ongoing yet unmonitored, event. Three types of triggers are observed in the cases above: periodic, stochastic, and tacit. Periodic and tacit triggers are proactive; stochastic triggers are reactive.

Periodic triggers are proactively clock-driven. Both the particulate matter and flood risk cases are instances of proactive, clock-driven triggers. In both cases, monitoring of system processes for events with significance to system integrity, and the subsequent evaluation of new information occurs before (in preparation for) scheduled (periodic) rule evaluation. In terms of planning for adaptation, this is probably the most intuitive of the triggering mechanisms. It is assumed that the system is continuing to change, but not at a rate that warrants an active rule adaptation scheme.

McCray et al. (2010) argue that “[t]he EPA program for ambient air standards is likely the most fully developed existing program of Planned Adaptation”. Stability of rule systems is the primary objective of conventional bureaucracies. Periodic triggers are arguably the most stable form of adaptation. Periodic triggers are predictable and fit nicely into an organization’s budgeting cycle. Periodic triggers allow managers to set aside resources (invest in capacity) for an upcoming evaluation, and provide technical analysts with the lead time to perform the kinds of rigorous, systematic analyses they prefer. The evolution of flood risk management, culminating in the Water Defense Act in 1995 and the Adaptive Delta Management framework in 2010, is an instance of learning from ad hoc adaptation to develop a more stable model of (periodic) adaptation.

Stochastic triggers are the direct product of active, continuous monitoring of a system for events that may impact (or have already impacted) the integrity of the system. Unlike known events monitored under a periodic triggering model, the stochastic model is monitoring for classes of events that may occur at any time. The stochastic trigger is a signal that an event has occurred, and that information from that event should be evaluated. Among the cases presented here, airline safety is the canonical instance of a stochastic trigger. Accidents are the events that are monitored by the NTSB. When an accident occurs, an investigation ensues. In this case, a stochastic event occurs, monitoring apparatus recognizes the event and invokes resources necessary to first cope with the impact of the event (mitigation), but, ideally, with appropriate investment in capacity, in short order, to process the implications of these events for the rule system in order to remediate.

Stochastic triggers pose a problem for a conventional organization’s capability and capacity planning. In contrast to periodic triggers, the investment is in monitoring the system, and the readiness to act on events that may not happen. Rules of recognition provide an authoritative mapping of event types and the bodies that should evaluate

those types. Rules of adjudication document the process for evaluation or, in some cases, in conjunction with rules of recognition, delineate the scope of discretion of the evaluating body. Rules of recognition here are, as may be obvious, not static, but are constantly updated, by rules of change *based on* the evaluative process. With the exception of very well-understood systems, the intersection of rules of recognition and adjudication will be concerned with the *scope* of expert discretion (recognition) and the distributional consequences of evaluation under that discretion (adjudication). This is especially important for high clockspeed stochastic events, where a series of independent evaluations, potentially by different actors, may lead to inconsistent recommendations. Trigger frequency is an important factor. In the simple case, one may assume that trigger frequency is one to one with event frequency. In the case of high clockspeed stochastic events, an aggregation rule may need to be in place to make the evaluative process more efficient and avoid unintended inconsistencies or worse yet, conflicting recommendations. The character of the evaluation process will be discussed shortly.

Tacit triggers are unique in that the trigger is based on a preliminary evaluation of the event. In this chapter, tacit triggers drive policy development in the RIR system. A policy change in the RIR system requires a rationale (a preliminary evaluation) before active consensus is invoked. In this case, the event is a change in industry practice, the numbers market, or the demand for numbers that warrants a change in policy. Tacit triggers are not monitored by the organization (the RIR and PDP shepherds) but rather by the affected stakeholders. In this case, the preliminary evaluation, not the event itself, triggers the process for changing policy. Hence, the trigger is tacit. Evaluating the tacit character of triggers in numbers policy is subtle: as noted in the case on RIRs, the ultimate source of an event is exogenous, but the recognition of the event is endogenous; adaptation is triggered internally by those experiencing the effects of those changes. From outside the system, because tacit triggers are driven by internal events, they may seem stochastic. The next subsection distinguishes among endogenous, exogenous, and boundary events.

In the abstract, tacit triggering is a process in which those actors bound by the rules have the option to propose change at any time. Returning to the issue of the clockspeed of triggers, the option to change policy is continuously available, but the organization has the option to provide feedback on feasibility, cost, and implementation timelines. This review process is also part of the adjudicative process, although, like in the RIRs, it may be authoritatively assigned to a different group of experts than those that evaluate the substance of the event itself. Tacit triggers, like stochastic triggers, may occur in rapid succession. In contrast to stochastic triggers, especially those driven by safety, tacit triggers do not have the immediacy of stochastic events. Tacit triggers still require an investment in on-demand capacity to deploy evaluative capabilities, but there is more flexibility in the deployment of those resources, especially within the timeline of the PDP.

Events

Events have been alluded to in the discussion of triggers, in terms of whether the trigger is a consequence of an event or an evaluation. Another important conceptual

distinction is whether the cause of the event is endogenous to the system, exogenous to the system, or whether it is a product of processes at the system boundary. These are referred to as endogenous, exogenous, or boundary events, respectively. In the cases discussed here, events are primarily labeled as belonging to one of these three categories, but this does not mean that systems cannot be affected by multiple categories.

The notions of endogenous and exogenous events are ideal types; no instance fits their criteria perfectly, but the rulemaking process may treat the event as one or the other as a way to scope secondary rules. *Endogenous events* are scoped to the system at hand: the system is shaped almost exclusively by actors managing the system, and the effects of the events are largely experienced by those actors. *Exogenous events* are scoped to the context in which the system being managed (and adapted) operates: these events affect the performance of the rules governing the system, but are not necessarily under the control of actors involved in governing either the system or the rulemaking process. *Boundary events* combine the two. Boundary events are exogenous in the sense that the cause is outside the scope of the rules, but those *causes* may be within the scope of influence of the rule system.

Numbers policy adaptation in the Internet case is largely driven by endogenous events. Application layer changes do affect number utilization, but tacit triggers activate an evaluation process that updates delegation rules to adapt to new conditions. The event is the internal consequence of a mismatch in policy with utilization demands. The effects of that event are borne largely by the numbers community. That said, the numbers system, whose function is defined by the rules dictating how it is managed, is within the control of the stakeholder collective affected by those rules. In this sense, both the events and the tacit triggers are endogenous to that system.

In the case of flood risk and airline safety, the events are exogenous. The events are driven by external factors that vary in the degree to which changes in the rules can control the manifestation of the event. In the case of flood risk, the event is typically a natural event, which is out the control of the rulemaking system. Policy adaptation is a process of monitoring natural events and developing mechanisms for countering the effects of these events. The event is not eliminated, as in the case of the numbers system, but the effects are mitigated. In the case of airline safety, systemic flaws are triggered by exogenous events such as weather, unexpected system conditions, or design flaws. Again, none of these events can be controlled or predicted, but the effects can be mitigated by evaluating the impact of a previously unobserved event and adapting rules appropriately. Note that while the particular accidents (events) are typically previously unobserved, the classes of events are not. The research efforts noted in the NTSB case are an effort to better characterize the classes of these exogenous events and identify patterns.

The particulate matter case is an instance of coping with long-running boundary events that are evaluated periodically. Particulate matter emissions are governed by EU and EPA standards in Europe and the US, respectively. That said, exogenous factors, such as innovations in industry, changes in industrial processes, and changes in the factors affecting industry outputs can all challenge the efficacy of EU and EPA rules. In that sense, the interaction among these exogenous variables and the

rules governing particulate matter levels occur at the boundary. Experts evaluate these events and make recommendations to the EU or EPA to update rules that will mitigate if the source of the event is out of the scope of regulation. If, on the other hand, the event is within the scope of standards, such as industry emission levels, remediation makes endogenous some or all of the root causes of the event itself.

For each of these classes of events, the role of evaluation is to understand the characteristics of these events and how susceptible they are to the system's rules. Evaluation is discussed in the next subsection.

13.3.3 Evaluation

Evaluation takes place at various points in the rule adaptation process. Organizational factors affecting evaluation are characterized in terms of the timing of evaluation, the composition of the evaluating body, and how coupled that body is with the actors adversely affected by inefficient rules.

Timing of Evaluation

Timing of the evaluation has been alluded to in terms of both the type of event and, more specifically, the type of trigger. The timing of the evaluation refers to both when the evaluation is activated and the duration of the evaluation. Evaluation may comprise monitoring and/or investigative capabilities. This discussion will distinguish two classes of evaluation: active assessment capabilities and vested assessment capabilities. Like the characterization of events, these are ideal types. Some cases may combine elements of active and vested assessment capabilities.

Active assessment capabilities are those that continuously monitor status of an event (or class of events), collecting information necessary for expert knowledge assessment of the state-of-the event. Active assessment is applied to events that are known to be in progress, but whose effects have yet to manifest. Flood risk management is the canonical instance among the cases presented here. Water defense is an ongoing aggregate event whose status is actively and continuously monitored to mitigate adverse effects. Particulate matter is another instance of active assessment. PM levels and effects are continuously monitored.

Vested assessment capabilities are an investment in the resources (capacity) necessary to respond to unanticipated exogenous events. Vested assessment capabilities comprise the investigative and analytic capabilities necessary to evaluate the effects of an unanticipated event. Vested capabilities are only initiated when an event occurs. Ideally, vested assessment capabilities can quickly offer a means to first mitigate the adverse effects of an event, then formulate a remediation if the effects of the event are within the scope of influence of the rule system. Airline safety is an instance of vested assessment capabilities. The NTSB "go-teams" are an investment in investigative capacity that can be deployed quickly after an accident (event).

The consensus process in numbers delegation policy development comprises elements of both active and vested assessment capabilities. Evaluation is active, in the

sense that those that experience inefficiencies in the rules are also those that create and update the rules. In the course of their operations, they monitor the efficacy of the rules. It is known and expected that patterns of delegation will change, but it is unclear precisely when. In this sense, events are unanticipated. The capacity necessary to support the consensus process is a vested assessment capability that can be activated by a member of the numbers community when inefficiencies cross thresholds agreed upon by the community. Although the trigger is tacit, as opposed to the stochastic trigger in the airline case, it signals an event that warrants evaluation to determine if adaptation is necessary. Together these capabilities signal an event should be considered, evaluates the effects of the event, and provides a forum for proposing solutions as a means to remediate the problem by updating delegation rules.

Loci of Assessment

Assessment processes may be concentrated within the same organization that makes the rules, distributed across one or more organizations, or (as in the case of numbers policy) distributed within an epistemic community. The *degree of concentration* is a function of the observability of events, and by whom. In the ideal concentrated form, the rulemaking system has the event observation and assessment capabilities and capacity necessary to generate sufficient triggers and perform event evaluation on its own. In the first distributed form, adaptive capabilities are located in an organization or distributed across organizations, separate from the organization making the rules. In the second distributed form, adaptive capabilities are distributed among individual experts that contribute to monitoring and evaluation. These ideal forms create a spectrum of loci of assessment.

The ideal concentrated form is included as a reference point. Harkening back to Hart's analysis, in historical governmental rulemaking systems, especially legal systems, the knowledge necessary to update rules is part of the knowledge base related to prosecutorial and judicial knowledge, rooted in that epistemic community. In the cases presented here, nearly all of the adaptive processes are distributed to some degree, requiring the rulemaking process to consume knowledge from a separate organization generating that knowledge. Two of the cases have a historical concentrated form: flood risk and the NTSB. Ad hoc adaptation in the Rijkswaterstaat was conducted by its own civil engineers. In the case of airline safety, the NTSB was initially managed by the FAA. In both cases, planned adaptation now follows a distributed form.

Of the cases as they stand now, among these distributed forms, the case of the NTSB is the most concentrated. In this case, there are two organizations: the FAA and the NTSB. The NTSB comprises experts in aviation operations and safety. As indicated by McCray et al. (2010), the NTSB occasionally enlists experts from the FAA. In this sense, along with the NTSB history, the NTSB is independent of the FAA, but its expertise is in certain cases coupled with the FAA.

The next most concentrated of these adaptive systems is the case of Adaptive Delta Management. In this instance, active assessment is performed by a number of expert organizations, but is funded by the government. This distributed form comprises a

single organization consuming new information and knowledge, but produced by a diverse set of monitoring and research organizations. Under this model, there is the potential for contention over the findings of these distributed monitors and assessors. One interpretation of this contention would be that it confounds creating stable, consistent rules. Another interpretation is that contention facilitates more effective analysis of risk and uncertainties. Contention highlights precisely where leading knowledge and analysis of events conflict—in effect, where uncertainties in the understanding of the event and its implications remain. Rather than rely on a single, potentially flawed analysis, this contention highlights where rules should incorporate contingencies to cope with unanticipated consequences, or at least to know which events it can mitigate and/or remediate, and which will require more ad hoc measures.

The standards for developing particulate matter are even more distributed. While the effects of PM are relatively localized, both rulemaking and the assessment capabilities necessary to inform these rulemaking processes are distributed. These two standards development processes draw on common bodies of (contending) scientific knowledge. In this case, there is variance in the kinds of triggers that drive adaptation, how evaluations are interpreted, and the evaluations themselves. Petersen et al. (2006) indicate that the EU's lag in adopting PM_{2.5} standards is explained in part by what has been characterized here as the variance in a distributed locus of assessment.

Lastly, number policy adaptation is the most distributed, but is scoped to a rather mutable system. The body of evaluators comprises individuals from a variety of firms that depend on IP addresses for their value proposition. In this sense, the loci of assessment are highly distributed. That said, recall that events are endogenous and the system is in effect a set of delegation rules and the supporting infrastructure for managing this delegation. As such, while highly distributed, the events that create demand for adaptation are well within the scope of the rulemaking process. The result is a highly decentralized evaluative body, but one that has all the necessary organizational levers to mitigate and remediate largely endogenous events.

Coupling

Closely related to, but conceptually independent of the loci of assessment is how coupled the loci of assessment are to rulemakers and those affected by events. Monitors and knowledge assessors range from tightly coupled to loosely coupled. Coupling has two factors: overlap and influence. *Overlap* means that some subset of the actors performing assessment is also the actors responsible for updating rules. *Influence* means that some subset of the actors performing assessment is under the influence of actors making the rules, typically because of a funding or employment relationship. Tightly coupled means that there is substantive overlap in the actors performing assessments and/or there is substantive influence over assessors by the organization responsible for rulemaking. Loose coupling is the other end of the spectrum: there is little to no overlap and/or influence.

Among the cases presented here, numbers policy is the most tightly coupled. Under the consensus process, there is a high degree of overlap: the assessors are effectively the rulemakers. That said, in most of the RIRs, the evaluation of the process is performed by the board or executive committee of the RIR. This introduces

a check on the discretion of the assessors and policy shepherds. While the numbers case is tightly coupled in terms of overlap, it is loosely coupled in terms of influence. The default, tacit mode of influence is that the numbers community must comply with numbers policy. That said, numbers community members are loosely coupled because they are employed by firms independent of the rulemaker.

The cases of flood risk and particulate matter have a similarly structured coupling. Both draw assessments from actors engaged in scientific research on the corresponding events. In terms of influence, there is greater influence in flood risk given that monitors are (partly) employed by the Rijkswaterstaat, supplemented by researchers funded (partly) by the Delta Programme. In the particulate matter case, monitors are similarly employed by US and EU agencies, but there is much more diversely funded research on particulate matter in both the USA and the EU. As such, there is similarly moderate coupling of the monitoring component of assessment, but the influence over particulate matter research is lower in particulate matter than in flood risk. In both cases, the relative magnitude of influence is based on two factors: first, the diversity of inputs and sources of scientific research; second, what proportion of this is funded by the rulemaking organization. As will be discussed in the next section, if reduced coupling is perceived to improve the quality of assessments and attendant adaptations, influence can be reduced by incorporating knowledge sources that are funded by some other authority, or (ideally) funded independently of any political or rulemaking organization.

Of these cases, the least coupled is the NTSB. The NTSB is also an instance of decoupling to improve the quality of the assessment process. Recall that the NTSB was at one time tightly coupled with the FAA, but, after public backlash over an instance of political influence, the Congress made the NTSB independent of the FAA. Although now loosely coupled in terms of influence, there is some overlap coupling. FAA employees are occasionally incorporated into NTSB investigations, but the NTSB reduces influence by ensuring that the analysis and recommendations are performed by the NTSB itself.

13.4 Conclusions and Ongoing Work

This chapter has reviewed four instances of planned adaptation representing diverse institutional make-up: state-agencies, international organizations, and a private transnational regime. While instances of planned adaptation are rare overall, within the four cases reviewed there is substantive diversity in the mechanisms that animate monitoring, knowledge assessment, triggers, events, and the organizational context in which these processes play out. In this concluding section, the range of configurations that give rise to adaptive capabilities and capacity are compared. This range of configurations highlights observed configurations, configurations that are considered feasible, and those that are considered infeasible. Given this range, the implications for planning and designing for adaptation are presented. The chapter concludes with implications for future studies—in particular, the value of identifying issues that may

be ripe for a planned adaptation approach, and the value of analyzing instances of ad hoc adaptation (as the foundations of planned adaptation) and failed adaptation.

13.4.1 Combinations of Adaptive Capabilities

To illustrate the adaptive capabilities space delineated by the cases, Table 13.1 summarizes the characterizations from Sect. 13.3.

A number of these characteristics are interdependent. Across the state-based cases, there is not an instance of tacit triggers. That said, as noted earlier, in conventional rulemaking, as articulated by Hart, the locus of adaptation is the state. In the conceptual vernacular established here, the closest to this form is the numbers community, a private transnational regime. In cases in which a government agency is the rule-maker, adaptation is driven by exogenous or boundary events—there is little overlap with the evaluator. This seems to indicate that these cases illustrate the need for some decoupling of monitoring, knowledge creation, and knowledge assessment from a potentially politicized rulemaking organization. In the case of PM and flood risk, the evaluator is a combination of the state and scientific research. The NTSB is an interesting case because, while the evaluator is loosely coupled, it is another agent of the state. Numbers policy is the counterexample in terms of coupling, but is substantively different because the rulemakers and evaluators are effectively the same body.

At one extreme, the government instances could be explained by simple economic specialization. Rulemaking organizations need some fluency in the system they are regulating, but do not necessarily have direct day-to-day experience with the management of that system. The job of the rulemaking organization is moni-

Table 13.1 Adaptive capabilities configurations^a

Aspect/Case→	PM	Flood risk	NTSB	Numbers
Organization	State, IGO	State + region	State	Private transnational
Trigger	Periodic	Periodic	Stochastic	Tacit
Event	Boundary	Exogenous	Exogenous	Endogenous
Clockspeed	4–7 years	6, 12 years	1–6 months	6–18 months
Evaluator	State, Research	State, Research	State	Industry
Timing	Active	Active	Vested	Active + Vested
Concentration	3	2	1	4
Overlap	Loose	Loose	Loose	Tight
Influence	Loose	Moderate	Loose	Loose

^a**Timing** refers to assessment timing; **Overlap** and **Influence** refer to types of coupling; **Concentration** refers to the concentration in the loci of assessment, following the relative ranking described earlier, where 1 is the most concentrated and 4 the least

toring and enforcing rules, which is conceptually different from monitoring events and designing triggers. This view explains the scarcity of planned adaptation among government agencies, but is a bit too simple.

A better explanation is that these agencies' coupling represents variants of the knowledge-policy interface. Among the state-based instances, two are combinations of state-based monitors, whose data are evaluated by both the state and external scientific research organizations. Further, the two that combine state and research have moderate concentration relative to the high concentration in the NTSB and the extremely low concentration in the numbers case. The character of the trigger and the event offer additional insights. In both cases, events are known, low clockspeed events that are monitored and evaluated over longer timescales (order of years) than numbers policy and aviation accidents (order of months). Neither presents an immediately observable hazard event; rather, these events represent changes that culminate in an adverse effect.

In both cases in which the evaluator is a combination of the state and research, the trigger is periodic. Monitoring data from long clockspeed events are fed to external researchers who can, within the bounds of the review schedule, perform evaluations as they see fit. The potential trade-off is in when friction in the knowledge-policy interface manifests. During the trigger period, the loose coupling may reduce friction. During the review period, friction in the knowledge-policy interface may increase. This friction will be revisited in the discussion of "Planning and Designing for Adaptation" (Sect. 13.4.2), and then again in the "Implications for Future Study" (Sect. 13.4.3).

In these cases, the timing of the evaluation is contingent on the clockspeed of the event, but not necessarily the source of the event. For both long clockspeed events, the timing is active. For the shorter clockspeed events, the timing has a vested element. In these latter events, evaluation is triggered by the event, focused on mitigation and remediation in both the case of the NTSB and numbers policy. Of particular interest here is the capability to invest in evaluative capacity by both the state and private actors. Conventionally, states are considered more far-sighted than industry when it comes to managing risk and uncertainty. As illustrated by these cases, they may not have the capabilities and capacity to act on that foresight, but they do have the social and political incentives. Here we see two counterexamples to the more conventionally intuitive cases of planned adaptation (PM and flood risk). In the NTSB, political factors have aligned to invest in "go-teams" necessary to respond to high clockspeed events. In the instance of numbers policy, endogenous economic incentives have led to investment in a consensus process that accounts for obligations to system integrity necessary for function and continued economic benefits.

13.4.2 Planning and Designing for Adaptation

The configurations highlighted thus far are those observed in the cases. One contribution of this conceptual model is the potential to reason about unobserved configu-

rations and under what conditions they may be feasible. While Table 13.1 illustrates clearly feasible configurations, it also hints at those *not* depicted, and that warrant discussion.

The ideal form of government-based adaptation that is tacit, endogenous, active, and vested, with high concentration and both tightly coupled overlap and influence noted earlier, is clearly not evident here. In terms of design and planning, it serves as an ideal form to start from. Which of these elements need to be relaxed to create a more feasible form? The closest instance here is numbers policy. While this instance is successful, it is a constrained case. Part of its success is that the stakeholders and the rulemakers are experts in a very narrowly scoped issue area. Further, unlike government-based rulemaking, which is often confounded by issue tying, the private transnational numbers regime has been historically decoupled from conventional political issues. It does not have to compete for resources with other issue areas, and it has the benefit of rulemakers who have a deep understanding of the implications of both primary and secondary rules. This is not necessarily the case in government-based rulemaking.

Another configuration that is not depicted in Table 13.1 is the case of high clock-speed, exogenous events with periodic triggers based on active assessment timing. Earlier, it was highlighted that high clockspeed feedback loops may create more harm than good. Even if it is assumed that rules of change can keep pace, high-speed adaptation may be ultimately detrimental for system integrity. As also noted earlier, a prized element of both rulemaking systems and the obligation to system integrity is stability.

For such a configuration the question becomes: are the effects of high clockspeed events sufficiently low impact to wait for periodic review? PM and flood risk, while not following this configuration, seem to make this choice. In the case of flood risk, the events driving active, continuous monitoring may be high impact (cost) events, such as “The Disaster” and the political implications of being “slapped in the face” again. In the case of PM, the high impact (cost) event is the cumulative effects. In both cases, cost-benefit analyses balanced these to select the active model over a vested, non-periodic model of adaptation. In the next subsection, this will be revisited in terms of how future studies can incorporate cost-benefits in terms of not only the substance of the policy analysis, but the costs of the attendant organizational structures.

Yet another configuration to consider (a variant of the ideal form) is the case in which the evaluator is the state, but the overlap and influence coupling is tight. In effect, this configuration is the original NTSB, before it was made independent of the FAA. Like the ideal form, this configuration invites issue tying. It may also be limited to domains in which the state has access to sufficient event data and experts necessary to evaluate the status and implications of a given event. Modulo issue tying, the NTSB under the FAA did have access to such data. Many technical systems are influenced more by natural events (flood risk) and industry innovation (PM, numbers). In these cases, an active assessment model requires events be observable by evaluators; PM and flood risk are such cases, but numbers is not. Further, the diversity of research perspectives demonstrated in PM and flood risk is also necessary, but missing.

The scarcity of successful instances of planned adaptation may require a planning and design strategy that not only builds on successful cases, but requires the close inspection of instances of ad hoc and failed adaptation. The instances of unobserved configurations mentioned above illustrate the analytic beginnings of such a modeling strategy. The conceptual model presented here offers an initial set of building blocks and the analytic framing to systematically speculate on unobserved, yet potential, configurations. That said, these are speculations. In the next (concluding) section, this approach will be framed as the foundation for a theory-driven sampling of instances of successful planned adaptation, ad hoc adaptation, and failed adaptation.

13.4.3 Implications for Future Study

The model presented here is a first pass at a framework for intentionally designing for planned adaptation. As noted in the introduction (Sect. 13.1), the substance of policy adaptation—domain-specific knowledge applied in conventional policy analyses—has a substantive literature dating back to the 1950s. This chapter contributes a complementary model that, building on the existing planned adaptation literature, considers policy adaptation in the context of the rulemaking system, and where elements of adaptation play out in the organizational complex that generates knowledge and applies it the process of policy adaptation. This concluding subsection outlines three research directions for the future study of planned adaptation:

1. A survey to identify not only additional cases of planned adaptation, but also borderline cases, instances of ad hoc adaptation, and failed attempts at planned adaptation;
2. Case studies and (where possible) interviews informing the processes that contributed to cases identified in the survey above, to refine the model presented here;
3. Integration of this model into conventional, substantive policy analyses.

An effective model is not only capable of expressing successful outcomes, but must also be capable of expressing instances of failures. More precisely still, it must be capable of expressing the criteria for success and the range of success and failure. The survey proposed would use the model presented here to develop a theoretically informed sample. Building on the unobserved configurations identified above, this effort would first further elaborate this initial set of unobserved configurations, then perform a survey of the literature on science and technology policy to identify cases that fit particular configurations or hint at unidentified configurations.

Given this body of “natural experiments,” the next step is to perform a deep dive into these cases to precisely elaborate the mechanics of how efforts at adaptation either succeeded or failed. These case analyses would characterize each case in terms of the model presented here. That said, recall that this model was noted in the introduction as the first iteration in contingent theory building. As instances of successes and failures are characterized, the conceptual model will itself be refined

to better represent both successes and failures. Pragmatically, given sufficient refinement, such a model would not only identify the mechanics of individual instances of adaptation, but the critical paths from one configuration to another. Even in this limited sample, the case histories for both flood risk and the NTSB provide early illustrations of such a critical path: flood risk followed a path from an ad hoc configuration to a planned configuration; the NTSB moved from a configuration with tightly coupled influence to a more independent configuration that preserved necessary overlap. Ultimately, this model would move from a planned adaptation space comprising relatively independent configuration points to a family of planned adaptation configurations, replete with observed critical paths, annotated with necessary resource endowments and investments, from one configuration to another.

Finally, this chapter recommends integrating this conceptual model into conventional policy analyses. Like the discussion of the unexpected costs of PM standards by Petersen et al. (2006), policy analyses usually incorporate cost-benefit analyses, but generally do not account for the organizational costs of capabilities and capacity to adapt. This model provides the foundations for systematically identifying the various loci of organizational costs. In particular, the variables describing loci of assessment and coupling can provide a starting point for understanding the cost of friction at the knowledge-policy interface. Such an integration would contribute to moving conventional policy analyses from mathematical models to more immediately applicable prescriptions that would appeal to both those concerned with obligations to system integrity and bureaucrats charged with evaluating the tractability of a solution. This latter may create its own feedback loop, reducing the friction at the knowledge-policy interface for expert policy analysts confronted by organizational barriers to effective implementation of planned adaptation models.

References

- ARIN. (2009). ARIN Policy Development Process. Retrieved from <https://www.arin.net/policy/pdp.html>.
- Bijker, W. E. (2007). American and Dutch coastal engineering: Differences in risk conception and differences in technological culture. *Social Studies of Science*, 37(1), 143–151.
- George, A. L., & Bennett, A. (2005). *Case studies and theory development in the social sciences*. Cambridge, MA: The MIT Press.
- Haas, P. M. (1992). Introduction: Epistemic communities and international policy coordination. *International Organization*, 46(1), 1–35.
- Hart, H. L. A. (1994). *The concept of law*. Oxford, UK: Oxford University Press.
- Keohane, R. O. (1984). *After hegemony: Cooperation and discord in the world political economy*. Princeton, New Jersey: Princeton University Press.
- McCray, L. E., Oye, K. A., & Petersen, A. C. (2010). Planned adaptation in risk regulation: An initial survey of US environmental, health, and safety regulation. *Technological Forecasting and Social Change*, 77(6), 951–959.
- Petersen, A. C., & Bloemen, P. (2015). Planned adaptation in design and testing of critical infrastructure: The case of flood safety in The Netherlands. In *International Symposium for Next Generation Infrastructure Conference Proceedings*, 30 September–1 October 2014. Schloss Laxenburg, Vienna, Austria: International Institute of Applied Systems Analysis (IIASA).

- Petersen, A. C., van der Sluijs, J., Tuinstra, W., & Martin, K. C. (2006). *Anticipation and adaptation in particulate matter policy: The European Union, the Netherlands, and United States*, 10–11 October 2006. Washington, DC: TransAtlantic Uncertainty Colloquium (TAUC) Workshop.
- Resnick, P. (2014). On Consensus and Humming in the IETF, RFC7282. <https://datatracker.ietf.org/doc/draft-resnick-on-consensus/>.
- Sowell, J. H. (2015). *Finding order in a contentious internet* (Ph.D. thesis). Cambridge, MA: Massachusetts Institute of Technology.

Dr. Jesse Sowell (Texas A&M University, Bush School of Government and Public Service) studies the capabilities and capacities of the transnational institutions that support Internet and cybersecurity operations. In particular, Dr. Sowell is interested in the role of consensus-based decision-making processes as a form of credible knowledge assessment necessary to adaptation in large, complex engineering systems. Dr. Sowell completed a Ph.D. in Technology, Management, and Policy at the Massachusetts Institute of Technology in 2015. From October 2016 to August 2018, Dr. Sowell was a Postdoctoral Cybersecurity Fellow at Stanford’s Center for International Security and Cooperation. In August 2018, Dr. Sowell joined the Department of International Affairs at Texas A&M’s Bush School of Government & Public Service.

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the chapter’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

