Policy Sci (2016) 49:351–372
DOI 10.1007/s11077-016-9258-9

Policy messes and their management

Emery Roe[1]

Published online: 15 July 2016

Abstract This paper presents a framework for better managing policy messes and draws implications for bad and good mess management in policy analysis and management. The framework has three foci: (1) the cognitive space in which policy messes develop, particularly in terms of gaps between macro-designers and micro-operators; (2) the unique domain of competence within that space where professionals manage the resulting messes by virtue of their skills in recognizing system-wide patterns, formulating locally specific contingency scenarios and translating both patterns and scenarios in highly reliable services; and (3) the ability of those mess and reliability professionals to be reliable in their domain and with these skills by maneuvering across different performance modes as conditions dictate—just-in-case, just-on-time, just-for-now or just-this-way.

Keywords Policy and management messes · Policy analysis · Public management · Wicked policy problems

Introduction

Start by taking a long, hard look around you: The world's a mess, and we know it. But if most everything is a mess, is each being managed for the mess it is? The answer is "No" for many policy and public management messes, and this greatly affects our management of the issues, now and into the future.

For the present, think of a policy mess as an issue so uncertain, complex, incomplete, or otherwise disputed that it can't be avoided. The question is: How to manage? The ideal is to prevent the mess, or clear it up and solve it once and for all, but that's easier said than done (in this article mess means specific messes). Yet every day, professionals reliably

✉ Emery Roe
emery.roe@berkeley.edu

[1] Center for Catastrophic Risk Management, University of California, Berkeley, Berkeley, CA, USA

manage to provide critical public–private services (outputs or outcomes), including water, electricity, transportation, telecommunications and even a good many financial services. They do this not by getting rid of their policy and management messes as much as by continuously sorting them, especially when those services are needed most—right now, in real time. What saves policy messes from being "wicked policy problems" or otherwise termed "unruly" or "intractable" is that mess is not all bad, some bad messes can be managed so as not to worsen, and good messes are from time to time to be pulled out of this.[1]

Policy messes are typically blamed on politics, dollars and incompetence. There's no uncertainty here as to the solution: If only we had better politics; if only money didn't drive the politics we have; if only those who controlled the money weren't incompetent, we wouldn't be in these messes. Even were this line of argument true as far as it goes, it doesn't go far enough for messes in the midst of high unpredictability that are the focus of this article.

Were messes no different than problems, we could rely on conventional policy analysis to get out of them. No such luck. Mess has long been an irreducible feature of my profession, policy analysis and public management, with its darkening sky of: coping, muddling through, groping along, suboptimization, bounded rationality, garbage can processes, second-best solutions, policy churn and policy fiascos, rotten compromises, managing the unexpected, crisis management, and that deep wellspring of miserablism, implementation and policy failure. Since these predators circle around the same prey, this article takes a closer look at the animal itself: the policy and management messes we find ourselves in when it comes to the reliable provision of important services like water, energy, transportation, telecommunications, health and the environment.

Of my two lines of argument, the first is the more important: Policymakers in government and policy analysts in the public and private sectors have a great deal to learn by way of management from a special class of professionals little discussed in the literature or the media: namely, those control room operators who manage critical infrastructures that societies have come to depend upon for reliable health, safety and energy services, to name a few. My book, *Making the Most of Mess: Reliability and Policy in Today's Management Challenges* (Roe 2013), is about applying what has been learned from managing more reliably in one domain (critical infrastructures) to the broader domains of policy and management that face their own political and legal mandates to be reliable, yet fall short of meeting the mandates.

Why would policymakers, analysts and political elites have anything to learn from control room operators? Because control room professionals manage to prevent all manner of major accidents and failures from happening, which would occur if not managed the way they are. We see politicians, policymakers and their advisors operating at their performance edges in turbulent environments; what we don't see is that critical infrastructure

[1] For more on wicked and unruly problems, see respectively Rittel and Webber (1973) and Ansell and Bartenberger (2016). While indebted to the literature on decision-making under uncertainty and in the face of complexity, I highlight that issue conflict and incompleteness are core to policy and management messes as well. To be clear, issues are complex when more numerous, varied, and interdependent than before. Issues are incomplete when efforts to address them are left interrupted, unfinished, or partially fulfilled. Issues are disputed when individuals take different positions on them because of their uncertainty, complexity, and incompleteness. As a first pass, think of policy messes as having all four properties. The recent work of Andy Stirling is especially illuminating as a starting point in thinking through these differences (Stirling 2010; Chapter 5 in Roe and Schulman 2016).

managers are doing the same in the face of turbulence, but more reliably by managing the ways they do.

My second line of argument is to answer: Just what is "managing the ways they do?" Control room operators are often brilliant mess managers, and what is blazingly obvious is we need better mess managers when it comes to what seem to be intractable policy issues. If asked why I term the latter, messes, my answer is that is precisely what they are called by those responsible for managing them. No metaphor or argument by analogy here. The healthcare mess, social security mess, financial mess, Eurozone mess, regulatory mess—those are the terms used by the public, analysts and policy elites to identify the issues and their special challenges. What is less recognized is that the same messes can be managed more reliably and professionally than the public and their policy establishments acknowledge. (Think of reliability as the safe and continuous provision of a critical service, even during—especially during—turbulent times.)

The image that public has of control rooms—functionaries in a darkened room undertaking command and control in front of computer screens and surrounded by walls of grid maps—captures none of the daily, minute-by-minute, adaptations required of operators to meet all kinds of contingencies that rise unexpectedly or uncontrollably and which have to be dealt with if the critical service is to be provided safely and continuously. These skills and management approach offer, I argue, a more effective template than do many policy analytical and decision-making perspectives. Three specific points are focused on in this article:

1. The first thing good mess managers show us is that *messes are managed, not cleaned up*. Too many people insist that the way to clear up policy messes is by reducing uncertainty, simplifying complexity, resolving conflict and completing unfinished business—in short, getting rid of turbulence. This advice, I argue, is why there are so many tangled muddles in policy and management. Far better human beings, incentives, policies, laws and politics than we currently have may indeed be required along the way, but if that is true, so is the fact that such insistence frequently makes for more messes, as we shall see.
2. What should we being doing instead? *We learn from those professionals whose job it is to manage mess all the time*. Managing well for them means they manage messes reliably or reliability messily. From them we also learn that *mess management requires at least three skills: pattern recognition, scenario formulation, and the ability to translate pattern and scenario into a reliable service when it really matters now*. These professional managers do not achieve reliability in safe and continuous critical service provision directly by designing broad systems to govern all discrete operations. To be reliable, professionals and the networks in which they operate interpret what system patterns mean for the locally specific scenarios they face now and in the next step ahead.
3. Why the need for translating pattern and scenario into what is required immediately by way of reliability? Because macro-designs, be they policies, principles, or laws, have to be modified in light of *both* the local features in the case at hand *and* the broader patterns that emerge across a run of individual operations. Both have to be accounted for to achieve reliable services. This sorting out process of recognizing system-wide patterns, formulating local scenarios and modifying scenarios in light of those patterns is complicated, but it is the core of mess management, a framework for which we turn to now.

Infrastructure control operators as reliable mess managers: a framework

How then do we actually better manage policy messes? Start with the studies that underscore the role of a unique cadre of professionals in managing critical services reliably (Roe and Schulman 2008, 2016).[2] The wider organizational literature to which my colleague, Paul Schulman, and I have been contributing tells us that the drive to highly reliable management in critical infrastructures can, for heuristic purposes, be sketched along two dimensions: (1) the type of knowledge used in activities to make system services reliable, and (2) the scope or focus of attention for those reliability activities. Reliability management is grounded in knowledge bases that range from experience, based on informal tacit understandings of the activities, to formal or representational knowledge, where abstract principles and deductive models are also core to understanding. Knowledge bases blend induction and deduction variously through the assembly of different arguments and scenarios related to reliability.

The scope of those managing for reliability ranges from a position that assumes reliability is an entire system output, encompassing many variables and elements, to a position that treats each case of reliability as a particular event with its own distinct properties or features. Typically, scope refers to the different scales, ranging from general to specific, that managers must take into account when reliability matters. Knowledge and scope define a cognitive space for managers, where reliability—the continuous and safe provision of the critical service even during turbulent periods and now not just in critical infrastructures—is to be pursued. Things start to get messy if only because perspectives of those operating within the space vary in terms of their knowledge bases and scope.

In this cognitive space, there are four nodal activities (Fig. 1), each position being a different mix of perspectives along the two continua:

The nodes—call these positions within the mess and reliability space the principal hubs for managing reliably—are macro-design, micro-operations, pattern recognition along with anticipation and contingency scenario formulation. We will see how and why mess and reliability professionals operate in the domain bound by the latter two hubs.

At the extreme of scope and formal principles is the macro-design approach to reliable critical services ("macro-design hub"). Design—whether in the form of policy, law, mission statement, or blueprint—asserts that formal deductive principles applied at the system-wide level govern a wide variety of critical processes for service provision. In this corner of the world, design is meant to cover the operation of an entire system, including every last case relevant to providing system services. At the other extreme in this cognitive space is reactive behavior in the face of real-time challenges at the micro-level ("micro-operations" hub). Here reliability depends on the immediate response of individual system operators working at the event level rather than on pre-existing designs at the system level.

Designers, however, cannot anticipate every eventuality. Worse, the more "complete" a logic of design principles aspires to be, the more likely its full set contains two or more principles contradicting each other, e.g., "we must not commit genocide … except when authorized by the nuclear doctrine of mutually assured destruction." On the other side, operator reactions are likely to give the operator too specific or partial a picture, losing sight of the forest for the burning trees in front of them. Micro-operations, in other words,

[2] I thank Paul R. Schulman for the basic framework and its formulation, though he is not responsible for my adaptations. This section's discussion of the framework builds on earlier research work in (Roe and Schulman 2008), which also reviews research by others on infrastructures as diverse as air traffic control, nuclear aircraft carriers and nuclear reactors.

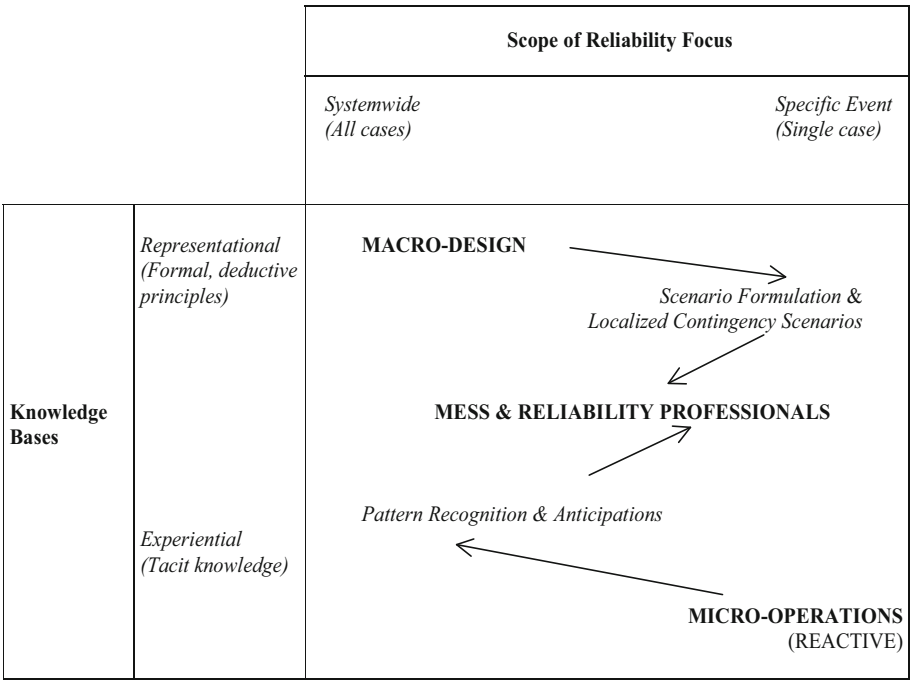

Fig. 1 Mess and reliability space for professionals (including policy analysts and public managers)

instills a "trained incapacity" undermining reliability, when operators are not aware of the wider ramifications of their activities.

What to do then, when high reliability is at stake? Moving across the cognitive space from one corner to its opposite is unlikely to be successful. Research finds that attempts to impose system-wide formal designs directly onto an individual event or case—to anticipate, fully deduce and determine behavior in each instance from macro-principles alone—are inadequate. On the other side, an individual's reactive operations scarcely constitute a tested template for scaling out to the system as a whole.

Instead of corner-to-corner movements, Fig. 1 indicates that reliability is enhanced when *multiple shifts in scope are accompanied by multiple shifts in the knowledge*. Becoming more reliable means becoming more knowledgeable about different things at different scales. Professionals approach and reach reliability through different skills than those for macro-design and micro-operations when it comes to managing reliability. Their approach is not direct, but indirect: messy in fact. Managers tack to reliability, much in the way that on a windy day a sailboat would not get from A to B via that proverbial single direct line.

We know from research that designers enhance reliability when they apply their designs less globally and relax their commitment to identifying principles that are meant to fully determine system operations. Both happen when designers contextualize design principles by embracing a wider range of contingencies in their analyses. They formulate alternate, more localized scenarios for system behavior and performance ("scenario formulation and localized contingency scenarios" hub in Fig. 1). As one from many examples, food policy

works much better when differentiating management protocols by type of crop or physical location (e.g., Godfray et al. 2010: 813).

We also know that reliability is enhanced when operations shift away from real-time reactions to recognizing patterns and anticipating their consequences across a run of cases of behavior and experience ("pattern recognition and anticipations" hub). Here "recognition" means looking for and into patterns, and "anticipation" means not only having expectations based on those patterns but also being prepared for their implications. Further, some patterns may be visible only at one scale rather than others.[3] By recognizing and anticipating patterns across cases, operators and managers learn to adapt, with better practices emerging across a run of cases. These anticipations and evolving strategies, based on empirical generalizations, trends, or other (quantitative or qualitative) patterns, are likely to be less formal than protocols developed through contingency analysis and scenario formulation. Note the term: better practices, not "best practice." A number of touted "best practices" confuse a scenario that works well in one case for the better practices, if any, that emerge across a run of cases and that then have to be customized, site-by-site, for effective results.

It is in this middle ground, bridging the formulation of design-based contingency scenarios to be realized more locally and the recognition of patterns and associated anticipations system-wide that we find the reliability-managing professional networked with like professionals. In the middle is where patterns and the anticipations based on them are probed, and where design-mediated scenarios are modified in light of the patterns at hand. It is in the middle where the skills in pattern recognition and locally specific scenario formulation reinforce each other; for example, repeated pattern recognition also helps build up skills in being sensitive to context-rich differences and vice versa. It is in the middle where reliability managers exercise skills of interpretation as pattern and scenario are translated into managing reliably. In the middle reliability manager must be mess manager: Here management is a craft and mess a term of art.[4]

Just who are these mess and reliability professionals? They are that unique class of professionals, generally middle-level managers and support staff, whose supervision, informal networks and skills ensure that critical services do not fail as often as they could. You find them in IT units, accounting units, engineering divisions, line operations, business continuity staff, inspectorates, auditing and budgeting departments, in some regulatory and legislative offices or on trading floors and in the field. Some senior executives, like a CFO or key staffer, might be part of the network, though scarcely "leading" all real-time operational decisions involved in managing reliably. Sometimes the middle-level staff and specialists appear in the press as self-identified "plumbers" (O'Connor 2008; Grant 2009: 21), but they rarely surface and when they do, even more rarely are the networks in which they work made visible in the process. In fact, some call their know-how "dark matter" (Hausmann and Sturzenegger 2005). That said, it is important to underscore that major executives cannot be expected to have this unique knowledge. Nor are mess and reliability managers to be confused with that lone but much-touted "street-level worker," "policy entrepreneur" and "change agent." The latter are better described in our framework as micro-operators or wannabe macro-designers. Where you see a mess and reliability

[3] In case it needs saying, pattern recognition and anticipation are not neutral or unmediated activities. We are witnessing species, be they protected whales and sharks or unprotected humans, equipped with transmitters so that they can be monitored in real time. Such efforts, however, individuate animals that often act collectively.

[4] My thanks to Martin Krieger for this latter point.

professional, you see many more networked to him or her, since no single individual could ever have sufficient pattern recognition and scenario formulation skills to operate in that unique domain on his or her own.

Bridging scenarios and patterns is the difficult part of the professionals' translation, since the interpolation involves transposing, transforming and synthesizing one in light of the other for present management. Little of this is mechanically step-by-step, e.g., first undertake pattern recognition and only thereafter shift to scenario formulation, because different professionals in the networked domain are skilled and interact in diverse ways. The translation of pattern and scenario into real-time action is interpretative rather than literal and is why new or different, a.k.a. unique, knowledge is generated, though not without its own difficulties.

We are now positioned to say much more than a policy mess is an amalgam of uncertainty, complexity, conflict and unfinished business and that it varies in terms of the performance conditions. A policy mess is any controversy or issue, the multiple and differing standpoints of which can be sorted out into the four hubs of macro-design, micro-operations, scenario formulation and pattern recognition. The task of management is to sort out the different positions at each hub and across hubs. Whether reliable mess managers in the networks of like professionals can extract a good mess from a bad one, or stop it from going bad, depends on their special knowledge to synthesize patterns and scenarios into reliable service provision. I see my own profession of policy analysis and public management is terms of just such networks of mess and reliability professionals.

It is tempting to see the hubs as a set of dualities, e.g., micro and macro, or system (patterns) and local (scenarios). But there are no simple dualities here because the differing knowledge bases ensure non-equivalence across the hubs. Note that in a policy mess, gaps in knowledge *always* exist between macro-design and pattern recognition (on the "system" side of the typology), between scenario formulation and micro-operations (on the typology's "local" side), and between pattern recognition and localized scenarios (the two very different sides of the domain of competence). When it comes to mess and reliability management, what one holds at the level of macro-theory for the system and what one finds in practice at the system level *necessarily* differ since each is based on different knowledge bases (so too macro and micro do not simply "mirror" each other). So too system patterns and the anticipations based on them are inescapably different from local scenarios that seek to contextualize design considerations better. That there is a gap, for example, between regulation on paper and as realized in a given region or compared to what is actually found in emerging practice across multiple regions scarcely surprises any serious professional. To act otherwise is to render management into bad mess management.

Bad mess management

The first task of a good mess manager is to avoid managing policy messes badly. There are many ways to screw things up in policy, but Fig. 1 highlights major ones around macro-design, scenario formulation, micro-operations and pattern recognition.

Going from bad to worst mess: no known patterns or scenarios

We saw how the autumn days of 2008 were fresh with panic after the bankruptcy of Lehman Brothers. "Frankly, everything is uncertain right now," a property developer told

a gathering of bankers, business leaders and decision-makers in Moscow. "We don't know whether to cut any contracts in roubles or dollars, or something else. We don't know what prices for anything will be, what demand will be, what our market will look like" (quoted in Tett 2009: 20). When the failure of practices hitherto grounded in trends and patterns combines with absent or otherwise wildly divergent contingency scenarios, nobody—and that includes the bridging mess and reliability professionals in banking and finance—can credibly claim to know what to do next.

It is one thing to be pushed into these unstudied conditions by earthquake, tsunami, or other disaster. It is quite another matter to rush wide-eyed into the unknown and then wonder why the rush turbo-charges panic and hyperturbulence. Turnover was so high in the financial trading sector before the financial meltdown that few remaining traders had experience with the derivatives-and-leverage collapse of the hedge fund, Long-Term Capital Management, in 1998. The chief executive of a New York-based hedge fund told the *Financial Times*, with only a little tongue in cheek,

> I think the best thing about working on Wall Street is that people can't remember what they did yesterday. It's remarkable. We do not use leverage in our distressed-security strategy. Now that times are better, some people have asked us, why don't you employ a little leverage? People forget. They said it wasn't so bad, maybe if we use a little leverage to enhance our returns, that's an OK thing to do. A little more time passes and maybe let's use a little more leverage, and before you know it we'll do it all over again. (Jeffrey Aronson in Freeland and Demos 2010: 16)

Nor are financiers alone in making such leaps. The president assures us about the presence of weapons of mass destruction that are not there. His second in command says they will be cheering in the streets when we arrive. His secretary of defense tells us the war will cost a fraction of what it actually cost. The secretary of state misinforms the United Nations, and the head of the major intelligence agency says the invasion will be a slam-dunk. And so on in what ended up as a conga line of ignorance because—and this is the point—when you are outside your domain of competence and intentionally in unstudied conditions, you can believe anything you want and ignore anything you don't want to hear. Other presidents and leaders have demonstrated this as well.

Once you are unable to calculate the probabilities and consequences of failure, all calls for better risk management are beside the point. You can't manage risks if they can't be estimated. So when reliability standards atrophied (as in housing appraisals leading up to the mortgage crisis) or became nonexistent (as in over-the-counter derivatives), risk is superseded by the incalculable. To be outside your domain of competence is to be undertaking activities that are dangerous precisely because "risk assessment and management" cannot navigate that unpredictability.

None of this is to insist that system patterns and local scenarios be clear in order to manage mess reliably or reliability messily. Policy messes are, to repeat, characterized by all manner of positions that are complex, uncertain, conflicted, or incomplete. That said, where no patterns or scenarios exist but where you want mandates for reliability in the driving seat, we should expect pressure over time to move from prevailing and conflicting macro- and micro-orientations to pattern recognition and scenario formulation. A non-financial example of a major management mess under such pressures is instructive.

For years, the use of fingerprinting, bullet, hair and handwriting analysis has revolved around a macro–micro-axis:

> Traditional forensic scientists seek to link crime scene evidence to a single person or object 'to the exclusion of all others in the world'.… They do this by leaning on the assumption of discernible uniqueness. According to this assumption, markings produced by different people or objects are observably different. Thus, when a pair of markings is not observably different, criminalists conclude that the marks were made by the same person or object. (Saks and Koehler 2005: 892)

The primary difficulty is this approach's reliability (Begley 2004: B1). A former editor of *Science*, Donald Kennedy (2003: 1625), long ago concluded,

> The problem… is that its reliability is unverified either by statistical models of fingerprint variation or by consistent data error rates. Nor does the problem with forensic methods end there. The use of hair samples in identification and the analysis of bullet markings exemplify the kind of 'scientific' evidence whose reliability may be exaggerated when presented to a jury.

In this case, the mess has always been in the details. A study of 86 wrongful convictions found that over 60 % had erroneous forensic science expert testimony as a contributing factor—the second most common one next to eyewitness errors (Saks and Koehler 2005: 893). Indeed, "error rates [have been] as high as 63 percent for voice ID, 40 percent for handwriting, 64 percent for bite marks, [and] 12 percent for hair" (Begley 2005: B1; see also Santos 2007: 4; Mnookin 2003). At these levels, dog sniffing is more reliable, with reported error rates for highly trained dogs of 30–40 % (David 2004: 42). As for eyewitness testimony, a US judge summed up: "Study after study revealed a troubling lack of reliability in eyewitness identifications.… Indeed, it is now widely known that eyewitness misidentification is the leading cause of wrongful convictions across the country" (quoted in Weiser 2011: A4).

From our framework perspective, it is not surprising that forensic science is being pushed to greater reliability by moving to the hubs of pattern recognition and localized scenarios via "developing measures of object attributes [for hair, fingerprints, teeth], collecting population data on frequencies of variations in those attributes, testing attribute interdependence, [and] calculating and explaining the probability that different objects share a common set of observable attributes" (Saks and Koehler 2005: 892). More reliable databases from which clearer patterns and practices emerge are a specific focus for "plugging" the holes in current forensic science (e.g., Fountain 2009). As "basic knowledge grows, experts will be able to inform courts about the relative strengths and weakness of their theories and methods, and suggest how that knowledge applies to individual cases" (Saks and Koehler 2005: 895). Whether or not this is a paradigm shift (e.g., Begley 2005: B1), it is very much a major movement to different hubs in the mess and reliability space.[5]

More types of bad mess management

There are other ways to be bad mess managers than running headlong into unstudied conditions. Here I focus on a few that are widespread in my reading, observation and work as a policy analyst: In terms of Fig. 1, decision-makers argue only from one hub or from

[5] In high reliability organizations, it is said professionals are as reliable as their last case. So too in this forensic science example. 'This is a business where you're as good as your last case', one of the past presidents of the American Academy of Forensic Sciences said about his profession (quoted in Hamill 2008: A12; for a book-length review of the issues, see Fisher 2008).

one standpoint at that hub; they confuse what is one hub for another; and they take shortcuts and bypass the unique knowledge and skills of mess and reliability professionals.

"The" hub, "the" standpoint

The mess here is that people treat an issue as if it were fixed and solved around one hub or a single standpoint. Getting the design right, the numbers right, the right person for the right job, or identifying the right scenario (not to mention asking the right questions) are the seductions of those who want to believe that messes can be cleaned up or avoided altogether. The appeal of starting and stopping with the micro-level individual or macro-level polity is so commonplace as to be nothing other than the origin and driver of many policy messes.

Confusion over hubs (with special attention to "prediction")

Decision-makers frequently mistake principles, patterns, scenarios and front-line experience for each other. The permutations are many, but I focus on a few that make for especially bad mess management.

Anticipations based on pattern recognition are often confused with specific contingency scenarios. Go back to the mid-2000s when Gary Becker (2005), Nobel economist, observed that the performance record of nuclear reactors was by and large positive. But his added point that this record constituted a rationale for going ahead and building a reactor somewhere specific is an altogether different matter. Anticipation that a trend will continue based on system-wide generalization is not a localized contingency scenario based on contextualizing a set of design principles. The scenario for *this* locality right now with *that* technology and *these* safeguards must first be posed and argued, irrespective of risk assessments grounded in frequency statistics across all operating reactors.[6]

One especially bad form of confusing scenario and pattern has been in post-9/11 risk assessment and management. It's fair enough that critical infrastructures and businesses plan and design for the "worst-case scenario" and formulate case-specific emergency protocols to ensure business continuity in the face of disaster. But that is the contingency scenario side of the professionals' domain. You cannot ask them to ignore the other side: the pattern recognition and the anticipations based on observed patterns across many business continuity efforts (Roe and Schulman 2008).

Macro-design and pattern recognition are also easily confused. The former head of Santander, Spain's largest bank, proffered advice for the 2008 financial mess: "Never buy a product you do not understand; don't sell a product you would never buy yourself; and if you don't know some of your customers extremely well, don't lend them money" (quoted in Betts 2008). That may well be sensible, but woe to those who take the advice as stable design principles instead of what may have been better-than-prevailing practices emerging out of dynamic conditions at that time.

Pattern recognition appears in each of the preceding confusions. This points to an added confusion: overlapping understandings of prediction. Remember, in our framework, pattern recognition is associated with anticipation; professionals base expectations and

[6] As we have seen, confusion can also go the other way when local scenario formulation is conflated with system-wide pattern recognition. A number of touted ''best practices'' confuse a scenario or protocol that works well in one case for the better practices that emerge across a run of cases and that then have to be customized, site by site.

preparations on the patterns observed, when it comes to trying to manage mess reliably. The problem is that the cognitive activity of anticipation is often conflated with prediction, while the other hubs also claim a role in prediction. The result is we have different phenomena passing for prediction without people realizing that they're actually talking about different things—which leads to more mess.

Clearly, macro-design can be seen as its own kind of prediction about what will happen if managers follow these principles rather than others. Localized contingency scenarios are also their own sort of prediction, when formulated as worst-case scenarios confronting management. Reactive micro-operations also involve prediction to the extent that the term "reactive" is based on what constitutes a professional response that follows the stimulus. Too often these different types of predictions are intermixed, when it might well be better to ask just which type of prediction is more at issue.

"Even the most intelligent and informed citizen (including lawyers and judges, for that matter) cannot predict with any reasonable assurance whether a wide range of seemingly ordinary activities might be regarded by federal prosecutors as felonies," argues a civil liberties lawyer (Silvergate 2009). In this example, are the laws ambiguous, which undermines their predictable application, or do the run of actual prosecutions not permit any firm prediction? Is it only that the federal courts have this problem when it comes to what are or are not felonies, or is it that even the best federal prosecutors react differently when it comes to making felony determinations? The mess with respect to felony assignments may not be that we can't predict, but that we're trying to predict all over the management space.

Shortcut leaps of faith

Much of managing policy messes poorly comes from decision-makers believing that reliability is achieved through leaps of faith from hub to hub that bypass the unique knowledge bases of mess and reliability professionals. Several cognitive shortcuts are significant because they are committed all the time.

Jumping directly from macro-design to micro-operations or the other way around Set the principle and everything should follow—or so we often hear. Someone asserts that each person has the same human rights as every other person. This move goes from a macro-design principle directly to micro-operations of personal identity. Those making this move are then upset when macro-principles, such as those in the UN's International Covenant on Economic, Social and Cultural Rights, are qualified by all manner of country-by-country reservations, understandings and declarations. The covenant guarantees our rights to education, marriage and holidays—except, that is, when declared in conflict with a country's constitution, laws, or religion.

From our framework, such reservations are *not* hypocritical but *must* be expected if human rights are to be reliably secured. It has been left up to nation states to enforce the universal values, and the only way we really know that human rights as macro-principles are taken seriously is to see how they are applied through local-specific scenarios, contingent country by country when not case by case. "Thou shall not kill' is all well and good, but we do not know how seriously that principle is treated until we get to grappling with qualifications such as 'except in cases of self-defense.' 'Granted that I should love my neighbour', wrote R. H. Tawney, British economic historian, but 'the questions which, under modern conditions of large-scale organization, remain for solution are, 'Who

precisely is my neighbour?' and, 'How exactly am I to make my love for him effective in practice?'" (in Caldwell 2008: 7). To ask "what is the law" is in effect to add the tacit suffix, "What is the law … in Roe v. Wade specifically" or "Hamdan v. Rumsfeld specifically" or in an equally specific case (MacCormick 2007: 5). (Note that just because we doubt human rights exist primarily as covering principles everywhere does not stop us from recognizing that we are at risk when systems behave as if better practices with respect to such rights were unimportant for case-specific situations in which we actually find ourselves.)

We as well see all manner of reverse micro-to-macro-metaphysics. Complex adaptive systems are said to arise autochthonously and nonlinearly out of micro-behavior (see Ehrlich and Levin 2005). It is also said that individuals acting under norms of economic rationality spontaneously generate efficient markets. Where is the mess here? Of course, individual micro-behavior can aggregate into patterns we can base empirical generalizations on. Yes, human greed led to patterned behavior that included financial bubbles; yes, policy design has a role in addressing them. Yes, bone and muscle have something to do with anatomy. But all this begs the question of who synthesizes the scramble of patterns and anticipations so as to manage the messes case by case.

It is clear that many policy messes we find ourselves in are due in no small part to the macro-to-micro- or micro-to-macro-approaches to governance. Get the principle right and democracy will follow; get individuals right and democracy will follow. Why wouldn't we conclude instead that really existing democratic governance is possible *only because* macro and micro are far less clear-cut across the policy and management messes we actually face? Yet privileging timeless principle and the individual over emerging practice and local contingencies in the face of undeniable turbulence is rampant in programs for public policy, economics, business and engineering.

Jumping directly from pattern recognition to macro-design and then to localized scenarios The move from the identifying systemic patterns to their supposedly direct implications for policy and legislation, and then posthaste to different program scenarios to "implement" that policy, is extremely popular. We teach policy students to do this all the time. Its popularity does not, however, make the move any less disabling when relying on it to manage our policy messes.

Global temperatures and the world's CO_2 emissions are increasing. *Therefore*, we must have a global strategy to deal with global climate change. That world-wide strategy *therefore* requires us to deal with, for example, the USA in one way and the Peoples Republic of China in another way as they account for most of the CO_2 emissions. There are, however, no *therefores* there.

What is missed in these leaps of faith is the unique knowledge base of mess and reliability professionals in the middle, whose task it is to make sense of the differing scenarios and patterns for the sake of dealing reliably with climate change, whatever its etiology. What is bypassed is how professionals are adapting to a climate change they do not know how to reduce or otherwise mitigate definitively (e.g., Adger et al. 2005). They include those who are continually searching for any better practices with respect to energy use, here and abroad, and who are very familiar with what it would take to translate and modify those practices so they would actually work for the case and context at hand. In fact, when it comes to global climate change, this manager is increasingly asked to take the region itself as the system of interest, not the globe, and then determine what regional better practices have to be modified in light of subregional or locality differences (Kennel

2009: 48–49). This is no incidental matter, as it is more plausible to imagine water and energy control rooms coordinating at the regional level than globally.

Just how does leaping directly from pattern recognition to macro-design end up making for more policy messes or making the ones we have worse? So what's really lost if you have to bypass mess managers from time to time, especially when "big picture" issues of design and global trends beckon? A great deal is lost and it is worth considering an example how this is so.

Consider three empirical generalizations that have currency today: (1) Megaprojects to construct huge infrastructures, such as dams and major roads, are habitually underestimated in terms of cost, overestimated in terms of benefits, and undervalued in terms of environmental impacts (e.g., Flyvbjerg et al. 2003); (2) a project developed incrementally is better than a megaproject implemented as planned (e.g., Easterly 2005); and (3) trade is often better than project aid (see Bhagwati 2002, 2005). Put aside for the moment the contrary evidence (e.g., Sharma 2005) and accept for the purposes of argument their status as generalizations over a wide distribution of individual observations.

Even if these statements were generalizable for the systems they purport to describe (as I believe they are), they scarcely justify jumping to macro-design principles that assert there should be no more planning for megaprojects; what projects there are should be smaller and incremental; and trade is to be preferred over project aid anyway. The three generalizations can in no way be taken to argue against localized scenarios that insist *under these contingencies* megaprojects and project aid are appropriate, as *here* other things are not equal. The burden of proof of course rests with those who argue for such specific-rich contingency scenarios in the face of system-level pattern recognition and anticipations to the contrary.

Other leaps of faith across the hubs The list of shortcuts can be extended easily. Here consider an example of an especially mess-inducing leap of faith, as exemplified by the controversy over the No Child Left Behind Act of 2001 in the USA. It is said that the application of this educational reform failed some regions and communities (a.k.a. failed local scenarios), *therefore* it failed children's education at individual schools there (a.k.a. failed micro-operations), and *therefore* we needed to redesign the NCLB Act or replace it altogether with some more efficacious macro-policy (e.g., Dillon 2005).

Yet there have long been more than 13,000 local school districts in the USA (Kraft and Furlong 2004). This very large number means there is likely a wide distribution of experience with respect to NCLB implementation—including that of other districts demographically like the ones that have failed. If so, why then would anyone recommend that we change the NCLB Act to make it work for failed sites and students without first identifying those districts that are similar to the ones that failed, but are doing better—even in the absence of new legislation (i.e., a new macro-design)? Wouldn't we first want to determine whether the failed sites could do as good if not better than their counterparts elsewhere under like conditions? Why would we ignore a bar that children, teachers or administrators have jumped, on the promise that overhauling macro-design will deliver the right bar everywhere? Avoiding the professional middle in such leaps erases the knowledge we already have and wastes resources, an inefficiency we can ill afford when the reliability stakes are as high as they are in US education.

Stopping messes from going bad or bad messes from getting worst may be the best, the very best, a decision-maker or policy manager can do in turbulent times. Many, however,

want more than prevention of the bad; we also want propagation of good mess management.

Good mess management

If bad mess management is insisting that we operate outside the domain of competence of mess and reliability professionals, then good mess management is what happens when operating inside the domain. Remember this domain, whose knowledge base differs from that at principle-based macro-policy and individual-based micro-operations, is bounded by the professionals' unique skills in system-wide pattern recognition and anticipation and in local contingency scenario formulation. But what specifically is going on in that domain of professional competence and by way of good mess management?

To telegraph ahead, the key to good mess management is the ability of networked professionals to maneuver across four performance modes in their unique domain of competence as conditions of task environment volatility and options change—all in the name of ensuring the continuous and safe provision of the service in question. There's the worsening mess of having to manage under just-for-now conditions (high volatility with few options) with firefighting and temporary fixes directed to keeping service provision underway. Better messes exist with just-in-case or just-on-time performance (high options whatever the volatility), though managing in either of these ways is not without its own risks. If there is a good mess in just-this-way management (reducing volatility through command and control), it is stopping a mess from becoming even worse. As for that best mess, staying reliable in the face of all the risks is being able to maneuver across performance modes as conditions change specifically with respect to volatility and options. From the other side, the worst mess is one in which it is not possible to work within any mode, let alone maneuver across them as conditions change. When that occurs you are coping, not managing, in unstudied conditions and in the precincts of bad mess management.

To see what the above summary entails, let's return to the control operators of large infrastructures who operate under mandates of high reliability. These professionals face a task environment of varying volatility, while the available responses to that volatility also vary. Volatility is the extent to which these system managers and operators confront uncontrollable or unpredictable conditions that threaten their ability to provide their system's critical service. Some periods are of low volatility: There are no surprising or unscheduled interruptions in the electricity supply, water provision, or financial services. Other periods are ones of high volatility: Temperatures go up, causing increased difficulties to the providers of electricity, water or health services. In some cases, volatility is high because what no one expected to happen actually does—for example, the Icelandic banking system collapses in a matter of days in 2008. Volatility, in other words, refers to the persisting or emerging instabilities in the task environment that confront the network or networks of managers, including system operators.

The managers we are talking about here have different resources in terms of money, personnel and strategies with which to address the volatility they face. This is called "option variety." High option variety means that an electric grid or investment firm has more resources available than the regulators require; low option variety means fewer resources are on hand to meet requirements. To reiterate, the systems we are talking about operate under reliability mandates. These may be de jure, as in the case of a bank's regulated capital reserve requirements, or de facto, as when a transmission operator

Table 1 Performance modes and chief risks for mess and reliability professionals

		Task Environment Volatility	
		High	*Low*
Option Variety	*High*	Just-on-time performance Risk: misjudgment with too many variables at play	Just-in-case performance Risk: complacency & inattention
	Low	Just-for-now performance Risk: amplification of deviance and error	Just-this-way performance Risk: lack of compliance

informally keeps a higher reserve of electricity than regulation mandates. Such reliability requirements can derive from system technology and/or organizational features. Supply and demand on the electricity grid (roughly, load and generation) must, for instance, be balanced to equal each other in real time, or the grid could collapse.

The two dimensions of task environment volatility (high and low) and options variety (high and low) set conditions for four performance modes that operators and managers work within as reliable service providers. The argument is that reliability in critical service provision (I am now thinking of critical services generally) requires access to all the performance modes, with each being its own form of mess management. While specific terms for the performance modes vary, for ease of reference I build on the terminology of our electricity research: "just-in-case," "just-on-time," "just-for-now" and "just-this-way" performance. Each is briefly described and then elaborated on as I discuss their features in terms of mess management (Table 1).

"Just-in-case" performance

When options are high and volatility low, many different options, resources and strategies exist "just in case" they are needed. Reserves are large, excess capacity exists and ample backups or fallbacks are available, all with little unpredictability or uncontrollability. This seems to be the ideal state of affairs in which to be (for the really existing manager, not necessarily for the economist), but it is not without its own messy risk. Operators and managers can grow complacent and end up not paying attention to changes in system volatility and/or options availability. Compared to the other modes, however, managing against complacency is a good mess to be in.

"Just-on-time" performance

When options and volatility are both high, just-on-time performance moves center stage. What worked today may not work under very similar conditions tomorrow. A specific

resource that was available just before could well not be available right now, and the manager has to be creative on the fly with the other options that remain. This performance condition requires real-time flexibility—that is, the ability to quickly make use of options, resources and strategies in order to meet the reliability requirements for safe and continuous service provision. Flexibility in real time means operators and managers are so focused in the moment on meeting a reliability requirement that they customize the match between the high volatility they face and the responses available. (This is why the performance should not be confused with just-in-time manufacturing: The latter can be just-plain-late when it actively discourages such flexibility.)

The major risk in just-on-time performance that combines creativity and discretion in how to sort out and assemble different options is misjudgment under the pressures of time and having too many balls in the air. Just-on- time performance means pulling a good mess out of ones that could go bad even in an instant.

"Just-for-now" performance

Using up resources can draw down the options available with which to respond, now and at the next steps ahead. When option variety is low but volatility remains high, just-for-now performance comes into play. "Just keep that valve open for now!" "Just stay late, that's all I'm asking!" "You've got to dial up the pressure from this point…" Just-for-now is the most unstable performance mode, and it is the one that operators and managers want to avoid most or exit from as soon as practicable. Why? Because they could well back themselves into a corner by trying to be reliable. In this mode, options and volatility can be linked (they are no longer independent dimensions), and being reliable now can make reliability all the more difficult to achieve later on. For example, operators and managers might have to go outside official channels or formal procedures to keep things reliable. Yet keeping equipment online when maintenance is overdue or insisting that already fatigued workers keep working longer can end up making things worse: Error, mistakes and deviance can be amplified which poses a major risk when there are few other options. Keeping something or someone working for just one hour more under these conditions may crash the system, even when one more hour would be nothing to worry other times.

From the standpoint of reliability, this performance mode cannot continue indefinitely. Operators and managers know they are not in control here; they know they are resorting to firefighting, band-aids and quick fixes. They understand how vulnerable the system is, how limited and interdependent options are, and they are busily engaged in trying to develop or secure resources to move out of this state. Just-for-now performance is such a worsening mess that, if protracted, it could become the worst imaginable—the system could fail entirely and end up in the realm of unstudied conditions so typically of bad mess management.

The prospect of prolonged just-for-now performance is one compelling reason why operators and managers are always on the lookout for better practices: What other such professionals do better in like corners may be a key way to increase options or reduce volatility in such worsening conditions.

"Just-this-way" performance

When the only option left is to reduce volatility directly, just-this-way performance moves front and center. One-way-only command and controls are asserted. A banking holiday is declared, mandatory job furloughs instituted, water conservation measures imposed and

shedding electrical load is enforced through scheduled blackouts. The great risk is that not everyone who needs to comply will comply, when following orders is the only way to ensure reliability. Just-this-way performance is stopping an already worsening mess from flipping into an altogether bad one, a full-fledged crisis.

Example of good management mess in the 2008 financial crisis

It is important to underscore that good as well as bad messes were witnessed throughout the financial mess—including events leading up to and after the 2008 crisis. Just-in-case management is what banks tried to do by increasing their capital reserves through bailout funds without, however, passing those funds on through increased lending; they held onto cash just in case things got worse. That was a good mess to be in from their viewpoint, but the way they managed their mess was part of the bad mess we borrowers found ourselves in at the same time.

Just-on-time management surfaced when liquidity was readily accessible when needed most. Liquidity in finance is the ability of a seller to assemble a deal when times get tough, which in our terms is the ability to assemble options, even if only at the last moment. Just-on- time liquidity is illustrated in the demise of the $6 billion hedge fund, Amaranth, in 2006, a collapse foreshadowing worse things to come. As Gillian Tett of the *Financial Times* described it then: "No sooner had Amaranth admitted to its losses, than buyers offered to purchase its gas portfolio (averting the prospect of dumping them on the open market)" (2006). This was a good mess in contrast to the 2008 bankruptcies that followed.

Just-this-way management in the form of command-and-control measures characterized a significant element of the meltdown. Those special government entities that were central to the US mortgage market, Fannie Mae and Freddie Mac, ended up nationalized. The UK government took over management of Northern Rock to stop a bank run. Direct command and control were asserted to ensure that mortgage rates and other lending became less mercurial—all in the name of securing greater authority over reducing volatility directly. An article in the *Financial Times* noted: "Tensions in money markets are so high we have witnessed the extraordinary spectacle of central banks not only providing liquidity but in effect becoming the market" (Davies 2008).

All of which leads to that most unstable performance mode, that clearly worse mess of just-for-now. The blowup of those novel, securitized financial instruments meant all manner of just-for-now relaxing of rules, special dispensations, emergency exemptions and one-time events like the FDIC's "systemic risk exemption" (on the latter, see Guha 2008). By way of example, in March 2008 the Federal Reserve arranged a last-ditch bailout for the nation's fifth largest investment bank, Bear Stearns, because the latter risked bankruptcy. How the bailout was managed is illuminating:

> In an action almost unprecedented in takeover history, JPMorgan bought 39.5 percent of Bear on the spot to ensure that it would have close to a majority of the votes to approve the deal. That agreement completely disregards New York Stock Exchange's rules that prevent anyone from buying more than 20 percent of [the] company without a shareholder vote. Other parts of the new agreement either stretch the rules or disregard years of precedent in Delaware, where both banks are incorporated. Of course, all of this rule-bending was done with the tacit, if not outright, approval of the federal government. (Sorkin 2008)

Paul Volcker, the former head of the Federal Reserve, summarized this just-for-now behavior as having taken the Fed to the "very edge of [its] lawful and implied powers" (quoted in Scholtes 2008). Paul Krugman, the Nobel Prize–winning economist, called it "barely legal" (2008). Other just-for-now transactions in banking occurred before the end of 2008, ranging from temporary lines of credit (Chan and McGinty 2010) to just-for-now circuit breakers and kill switches intended to "temporarily" interrupt market prices from falling below set limits.

Such distinctions between modes of mess and mess management matter greatly. First, proposed improvements all too often fall short of producing reliable financial services when effects of those proposals on options and volatility are not made explicit. Adding liquidity in a volatile financial environment means that the system can at best be managed just-on-time rather than just-for-now; in other words, increasing financial liquidity may add options but do little to reduce financial volatility. For instance, a central bank window for emergency lending or a treasury line of credit may be better at increasing options through added liquidity than at reducing financial volatility. Increasing capital reserve requirements or putting limits on leverage—the ratio of borrowings to equity—may, on the other hand, have everything to do with stabilizing financial volatility without necessarily adding more options for financial managers. A government plan for bank recapitalization to allay public fears over widespread bank insolvency may, in contrast, have just as much to do with reducing system volatility as increasing those banks' options. So clearing houses for over-the-counter derivatives may also have the virtue of addressing simultaneously the reduction in volatility and the increase in options, at least to the degree that the clearinghouses are not themselves so interconnected as to increase risks (Dudley 2012).

Let's now shift the discussion from good messes and worse messes to what were the best and the worst messes in the financial meltdown. As discussed in the section on bad mess management, the worst mess in the financial crisis—and here "crisis" is the correct term—was the panic recorded in the last quarter or so of 2008, when it was in no way evident just who was going to be saved and who would be left to fail. But where was the good mess—in fact the best mess to be in—during the financial crisis? To discuss the meltdown in terms of good messes seems counterintuitive or perverse. Yet the mess and reliability framework of this article suggests that the good messes are there to be found if analysts only looked.

For this author the best mess in the financial crisis was a very big dog that didn't bark in the panic. In conditions uniformly described as bad to terrible, the financial meltdown was accompanied by a positive and beneficial silence that went largely unreported. Succinctly put, while banking and finance are a global infrastructure, the financial mess (even when it morphed into a crisis of unstudied conditions) did not spread to other critical infrastructures as rapidly and pervasively as it did through the real economy—and that was a very good mess to be in.

The effect of the financial crisis on the real economy has been much noted and is not disputed here. In the words of the Financial Crisis Inquiry Commission: "Distress in one area of the financial markets led to failures in other areas by way of interconnections and vulnerabilities that bankers, government officials, and others had missed or dismissed" (2011: 27). True, but why didn't it spread further into other society's critical infrastructures? After all, critical infrastructures are said to be intensively coupled. "Interconnected systems too complex and dangerous to fail are not unique to financial services," writes the economist John Kay (2009). "Failure could also have catastrophic consequences in electricity networks, oil refineries and petrochemical plants." In an interconnected critical infrastructure world, where everything is commonly described as connected to everything

else, we would expect very real impacts of the financial crisis on other sectors such as transportation and natural gas as well, if only through the freezing up of credit, trade finance, and investment in shipping and public transportation (see Wright 2008, 2009).

The financial contagion certainly hampered infrastructural projects in transportation and ports (Little 2009). But what about that electricity network so core to many other infrastructures? Was electricity affected by the financial mess to the same extent as the real economy was? No, if we rely on those same reports about the financial meltdown. While the upheaval clearly affected plans for future infrastructure in electricity (see, for example, Crooks 2008; Wigglesworth et al. 2009), it did not affect real-time operations in any newsworthy way—save for one possibly major exception: the massive 2012 electrical blackout in India. That outage was connected to interrelated problems in India's electricity and banking sectors, though how much those problems were due to the global financial mess as distinct from political issues in India remains an open question (Sender and Crabtree 2012).

Since no one doubts that the real economy was hit hard by the financial upheaval, what is going on here? The short answer is that other infrastructures continued to manage, however messily, in the face of the financial turmoil and in ways demonstrably different from what was happening in parts of the banking and finance infrastructure. The only study I know of relevance to this issue that examines cross-infrastructure cascades has been undertaken by the Dutch research body, TNO Defence, Security and Safety, and the Delft University of Technology (Luiijf 2008). The TNO study concluded: "Our analysis of the collected data shows that most cascades originate from only a limited number of critical sectors (energy, telecom) and that interdependencies occur far less often than most theoretical studies assume" (Luiijf 2008). Few cascades were initiated by the financial services sector compared to the energy and telecom sectors. "Dependencies" among many infrastructures, write Eric Luiijf and his colleagues (2010: 16) in a later and longer review of the TNO data, "are anything but unmanaged." In short, our starting assumption should be we are in the world of professionals who manage all manner of messes more reliably than a good many others suppose.

Summary and conclusion

To summarize this article's framework, policy analysts and public managers are mess and reliability professionals charged with making sense of the wider patterns and locally specific contingency scenarios they face, case by case and often just on time. They are caught in the middle of things, where it's easier to say we're in a policy mess and here's how to manage it than it is to devise a narrative about how the mess began or will end. At their best, analysts and managers identify and sort out standpoints they confront around the four hubs of macro, micro, pattern and scenario, with an eye to rummaging the good mess out when it counts the most—for now and next step ahead.

When working under conditions that are highly unpredictable or uncontrollable, it is better to manage your options creatively up to the last moment (just on time) than manage reactively with fewer options over prolonged periods (just for now). One way to make conditions more predictable (just this way) is by declaring an emergency, but that course of action has built-in risks. Searching for better practices can well be the only viable avenue to reduce your task volatility and/or increase your options to respond to that volatility. That there are no guarantees in any of this—better practices may not be emerging or, if already

emerged, the practices may not be modifiable to the specific contingency scenarios being faced right now—should be absolutely clear by this point. More, all this is not without its own risks, but they are risks that are known or knowable rather than the dangers and hazards of working in altogether unstudied conditions. It bears repeating that to manage this way means you have to work within a network or networks of like professionals.

If the above is correct, then a great deal of rethinking needs to be done about what makes for effective policy analysis and management when it comes to services considered or mandated as "critical" (whether in an infrastructure or not). Allow me to conclude with three reconsiderations deriving from the important role of better practices in mess and reliability management (for more, see Roe 2013)

First, professionals in the middle domain of competence are more likely to ask, "Who"s figured it out better and how can that be modified for use here?' Local scenarios are melded with broader practices that have been found to work in similar situations, where the learning involved in modifying the practices to the specific site can and should feed back into the broader knowledge base of what works by way of management. If this sounds like old-fashion incrementalism—though it is nothing like the one I was taught—then spell it, Incrementalism, where the capital "I" indicates a scope of search for better practices that aspires to be international.

The key point to be made about the better practices that emerge out of a run of different cases is that *they import into management the scale of governance at which they actually work across that run of cases*. We can, of course, hope that better practices emerge across scales and levels of management, but to be clear, a practice that turns out to be better only across the same level of government or scale of collective action cannot be presumed to work at different levels or scales. Any such applicability is an empirical issue that cannot be settled a priori and to assume otherwise is to undertake bad mess management.

Second, it's necessary to rethink what constitutes risk analysis. Traditionally, the three big questions in risk assessment and management have been: What could go wrong? How likely is it?, and What are the consequences if it does go wrong? A risk analysis that takes policy and management messes seriously asks the same questions but in a very different order: What's going right? What's even better? How do we get there? and then: What could go wrong in trying to get there? How likely is that? What are the consequences if that goes wrong?

Third, one huge category error in conventional policy analysis and public management has been to assume that implementation is about converting macro-design directly into micro-operations and that implementation at the micro-level can end up as a kind of de facto policymaking for the macro-level. Actual implementation, however, is better understood as taking place in the middle across a network of professionals by those skilled in pattern recognition and formulating localized contingency scenarios. Here, the locus of implementation shifts away from micro-operators—the fabled street-level worker, including the cop on the beat, the teacher in the classroom, and the caseworker on a home visit (who may not even see themselves as "implementing policy")—to networks of skilled professionals and the risks they face. Here it is not the case that "the biggest problem we have is implementing policy," but rather: The biggest problem is to adapt better practices, where they exist, to policy and management issues that have to be reliable *beyond* macro-design and micro-reaction.

The implications of these three points, by way of conclusion, are profound for my profession, I believe. It is said we as a society have no choice but to experiment on a large scale in these turbulent times; we are, policy elites insist, in unprecedented conditions with no alternative but to undertake what was never done before. Necessity has no law

(*necessitas non habet legem*), as the Romans put it. But the Romans also had that other maxim: The welfare of the people is the supreme law (*salus populi suprema lex esto*) (Sonenscher 2007; Fatovic 2009). From the mess and reliability perspective, when it comes to governance we experiment only *after* we—repeat, *we*—have confirmed a number of others have no better practices in that regard.

Acknowledgments A shorter, different version of this article is to appear as Chapter 12 in C. Ansell, J. Trondal and M. Ogard (eds.) *Governance in Turbulent Times*. Oxford: Oxford University Press (forthcoming 2016). My thanks go to Duke University Press for permission to republish material from *Making the Most of Mess: Reliability and Policy in Today's Management Challenges* (Roe 2013). Because of space limitations, this article focuses on the core management approach, with readers referred to the book for detailed applications and references, including literature review and elaboration of the interaction of cultures, politics and organizations in mess and its management. In addition to those acknowledged in the book and the short chapter version, I wish to thank Paul R. Schulman for his contribution to the framework, Martin Krieger for his suggestions and encouragement, and the comments of two reviewers and the editor of *Policy Sciences*.

References

Adger, W. N., et al. (2005). Social-ecological resilience to coastal disaster. *Science, 309*, 1036–1039.
Ansell, C., & Bartenberger, M. (2016). Unruly problems, chapter 5. In C. Ansell, J. Trondal, & M. Ogard (Eds.), *Governance in turbulent times*. Oxford: Oxford University Press.
Becker, G. (2005). The nuclear option. *Wall Street Journal*, May 12.
Begley, S. (2004). Despite its reputation, fingerprint evidence isn't really infallible. *Wall Street Journal*, June 4.
Begley, S. (2005). Fingerprint matches come under more fire as potentially fallible. *Wall Street Journal*, October 7.
Betts, P. (2008). Europe's blameless banks fall off their pedestals. *Financial Times*, December 18.
Bhagwati, J. (2002). The poor's best hope. *Economist*, June 22, 24–25.
Bhagwati, J. (2005). A chance to lift the 'aid curse'. *Wall Street Journal*, March 22.
Caldwell, C. (2008) Time for some morality trades. *Financial Times*, December 20–21.
Chan, S., & McGinty, J. C. (2010). In crisis, fed opened vault wide for US and world, data shows. *New York Times*, December 2.
Crooks, E. (2008). Energy security will be hit by global slowdown. *Financial Times*, November 24.
David, G. (2004). Sit! Stay! Testify!. *Fortune*, January 26.
Davies, P. (2008). Real reason to regulate derivatives is liquidity. *Financial Times*, October 3.
Dillon, S. (2005). Report faults Bush initiative on education. *New York Times*, February 24.
Dudley, W. (2012). How we will stop derivatives magnifying future crises. *Financial Times*, April 16.
Easterly, W. (2005). Big plans to halt poverty don't work. *International Herald Tribune*, July 4.
Ehrlich, P., & Levin, S. (2005). The evolution of norms. *PLoS Biology, 3*(6), 943–948.
Fatovic, C. (2009). *Outside the law: Emergency and executive power*. Baltimore: Johns Hopkins University Press.
Financial Crisis Inquiry Commission (FCIC). (2011). *The financial crisis inquiry report: Final report of the national commission on the causes of the financial and economic crisis in the United States*. New York: Public Affairs Reports.
Fisher, J. (2008). *Forensics under fire: Are bad science and dueling experts corrupting criminal justice?*. New Brunswick, NJ: Rutgers University Press.
Flyvbjerg, B., Bruzelius, N., & Rothengatter, W. (2003). *Megaprojects and risk: An anatomy of ambition*. Cambridge: Cambridge University Press.
Fountain, H. (2009). Plugging holes in the science. *New York Times*, May 12.
Freeland, C., & Demos, T. (2010). View from the top: Jeffrey Aronson. *Financial Times*, March 1.
Godfray, H. C., et al. (2010). Food security: The challenge of feeding 9 billion people. *Science, 327*, 812–818.
Grant, J. (2009). Clearing not the cure-all for financial system woes. *Financial Times*, June 26.
Guha, K., et al. (2008). Rally defies deepening crisis. *Financial Times*, October 1.
Hamill, S. (2008). Pathologist accused of profiting from office. *New York Times*, January 28.
Hausmann, R., & Sturzenegger, F. (2005). 'Dark matter' makes the US deficit disappear. *Financial Times*, December 8.

Kay, J. (2009). ''Too big to fail' is too dumb to keep.' *Financial Times*, October 28.
Kennedy, D. (2003). Forensic science: Oxymoron? *Science, 302*, 1625.
Kennel, C. (2009). Climate change: Think globally, assess regionally, act locally. *Issues in Science and Technology*, *Winter*, pp. 47–52.
Kraft, M., & Furlong, S. (2004). *Public policy: Politics, analysis, and alternatives*. Washington: CQ Press.
Krugman, P. (2008). Success breeds failure. *New York Times*, May 5.
Little, R. G. (2009). A cautionary tale to return to investment basics. Letter to the editor. *Financial Times*, October 7.
Luiijf, E., et al. (2008). Empirical findings on critical infrastructure dependencies in Europe. The Hague: TNO Defence, Security and Safety and the Delft University of Technology. http://critis08.dia.uniroma3.it/pdf/CRITISe08e40.pdf.
Luiijf, E., et al. (2010). Empirical findings on European critical infrastructure dependencies. *International Journal of System of Systems Engineering, 2*(1), 3–18.
MacCormick, N. (2007). Mr. Justice: Ronald Dworkin's troubled search for respectful politics, a moral law and the fundamentals of democracy, TLS. *December 7*, pp. 3–5.
Mnookin, J. (2003). Fingerprints: Not a gold standard. *Issues in Science and Technology, 20*(1), 47–54.
O'Connor, S. (2008). Financial risk 'plumbers' hope for Corrigan boom. *Financial Times*, August 22.
Rittel, H., & Webber, M. (1973). Dilemmas in a general theory of planning. *Policy Sciences, 4*, 155–169.
Roe, E. (2013). *Making the most of mess: Reliability and policy in today's management challenges*. Durham, NC: Duke University Press.
Roe, E., & Schulman, P. (2008). *High reliability management: Operating on the edge*. Stanford, CA: Stanford University Press.
Roe, E., & Schulman, P. (2016). *Reliability and risk: The challenge of managing interconnected infrastructures*. Stanford, CA: Stanford University Press.
Saks, M., & Koehler, J. (2005). The coming paradigm shift in forensic identification science. *Science, 309*, 892–895.
Santos, F. (2007). Evidence from bite marks, it turns out, is not so elementary. *New York Times*, January 28.
Scholtes, S. (2008). Volcker urges extended role for the Fed. *Financial Times*, April 9.
Sender, H., & Crabtree, J. (2012). Powerless to act. *Financial Times*, August 6.
Sharma, S. (2005). He encouraging truth about foreign aid. *Challenge, 48*(4), 11–25.
Silvergate, H. (2009). The decline and fall of *mens rea*. Champion, September–October. Accessed online on April 20, 2016, at http://www.nacdl.org/champion.aspx?id=14738.
Sonenscher, M. (2007). *Before the deluge: Public debt, inequality, and the intellectual origins of the french revolution*. Princeton, NJ: Princeton University Press.
Sorkin, A. R. (2008). The hand behind the deal. *New York Times*, March 25.
Stirling, A. (2010). Keep it complex. *Nature, 468*, 1029–1031.
Tett, G. (2006). Why $6bn hedge fund loss left the finance world unruffled. *Financial Times*, September 22.
Tett, G. (2009). Financial world is stumbling blindly through a cognitive fog. *Financial Times*, March 6.
Weiser, B. (2011). In New Jersey, sweeping shift on witness IDs. *New York Times*, August 25.
Wigglesworth, R., Sakoui, A., & Kerr, S. (2009). Utility's $2bn debt woe raises pressure on Dubai. *Financial Times*, December 9.
Wright, R. (2008). Dead wait: Shipping. *Financial Times*, December 18.
Wright, R. (2009). Slide in trade leaves a glut of capacity. *Financial Times*, May 26.