

THREE

THE WIDER FRAMEWORK FOR MANAGING MESS RELIABLY

Hubs, Skills, and the Domain of Competence

It's rush hour and traffic is its usual mess. As luck would have it, three cars arrive simultaneously at the four-way stop, with traffic backed up behind each. The three cars' turn signals aren't on, but each driver can see the other two clearly. Drivers 1 and 3 are going in the opposite direction, with driver 2 to the right of driver 1 and driver 3 to the right of driver 2 (see figure 2). In the United States, the rules of the road are that the driver on the right proceeds first. On this principle, driver 3 begins to cross the intersection, and in theory driver 2 would go next, followed by driver 1. Driver 2, however, does not go next. Drivers know that in this kind of situation, in rush hour with cars stacking up, it is important to keep the traffic moving. So driver 1 crosses the road just as driver 3 does, and driver 2 goes afterward, none of the three seeing anything untoward in doing so. Of course, the three drivers have to keep alert in case something unexpected happens that requires an immediate reaction on their part.

All the principal features of managing a mess reliably are here. The rules of the road are design principles developed to make traffic flow reliably. Local circumstance can and does require their modification in practice, however. In these cases, reliability depends on the ability of the drivers to recognize then-pertaining patterns and formulate contingency scenarios. Skilled drivers familiar with a situation like the one described know that in these cases, moving across the intersection in the sequence these three drivers did can be quicker than keeping to the rule. In fact, keeping to the rule could make the traffic mess worse.

Even though our drivers may never have been at this intersection before at just this time of day, each formulates a scenario in light of the known patterns and proceeds on that basis. Those in cars farther back would be upset otherwise, as all three drivers are connected with others behind them in the ensuing traffic flow. Sometimes the mess goes from bad to worse regardless. If a driver gets a flat tire while turning

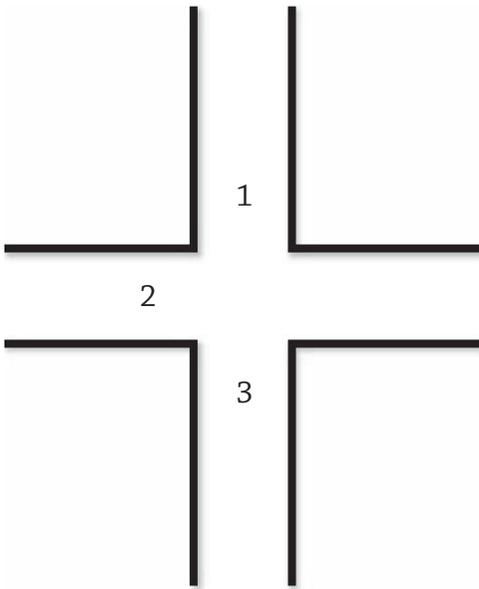


Figure 2. Managing a traffic mess

into the intersection, that may be a scenario none of the drivers present have witnessed before. A first-time driver in this area might get confused. Where no pattern or scenario exists, drivers have to depend on their reaction skills.

The basic elements of macro design, pattern recognition, scenario formulation, and reactive micro operations are roughly captured in this familiar example of managing in rush hour. The analytic framework for connecting these concepts was originally developed by my colleague, Paul Schulman, to understand how electricity control room operators think through their management under pressing time and knowledge constraints. Here and for the rest of the book, I move beyond control room managers and operators by adjusting the framework and its elements to describe reliable mess managers. In formal terms, I argue that these managers should be seen as “mess and reliability professionals” whose special skills, expertise, and competence lie in reliably sorting out the policy and management messes they confront and doing so within the politics they face.¹ I begin by introducing basic concepts and terms, starting again with service reliability, and show how the discussion relates specifically to mess. I apply the extended framework to what many consider the major messes driving economic and social upheavals—namely, overpopulation and globalization. Once the overall framework is presented, chapter 2’s fourfold typology of messes and mess management will be located within it. In

the process, I recast core concepts such as risk and learning, while showing how mess and reliability professionals differ from others. The chapter ends with a fuller definition of a policy mess and their difficulties. I defer discussion of the politics involved to chapter 6. Throughout, I talk about “you,” “we,” and “us” in the belief that we all are better mess managers in the making.

A General Framework

Leave the cars at the intersection and ratchet the analysis up to the system level. The wider organizational literature to which my colleagues and I have been contributing tells us that the drive to highly reliable management in critical infrastructures can be, for heuristic purposes, characterized along two dimensions:² (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, in which abstract principles and deductive models are also core to understanding activities. Knowledge bases blend induction and deduction in varying ways, which are reflected in the assembly of different arguments and scenarios with respect to reliability.³

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 get messy if only because the perspectives of those operating within the space vary in terms of their knowledge bases and scope. As we will see momentarily, things get even messier in the name of reliability.

In this cognitive space, there are four nodal activities (see figure 3), each position being a different mix of perspectives along the two continua. The nodes—I call these positions within the mess and reliability space the principal hubs for thinking about and managing reliably—

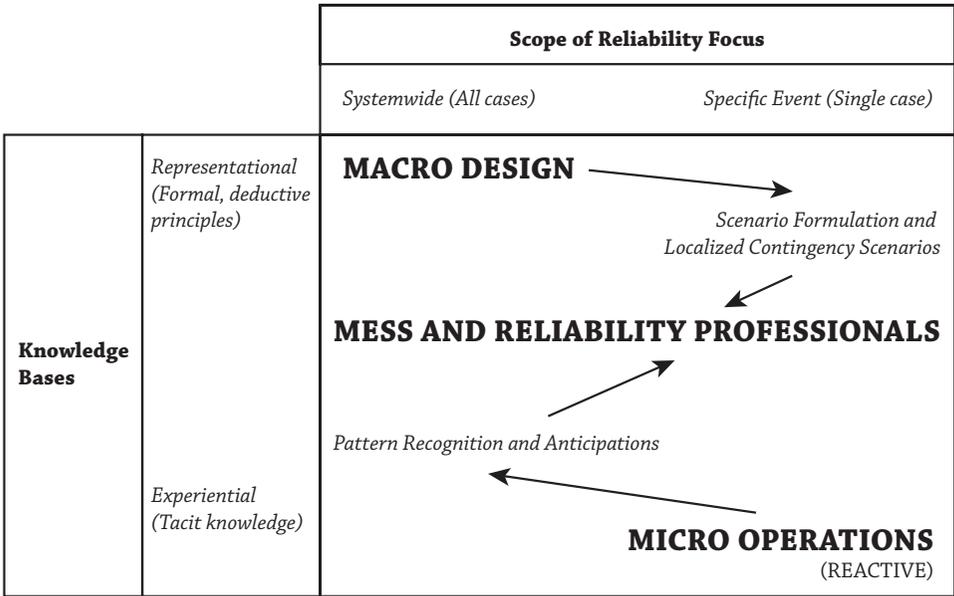


Figure 3. Mess and reliability space of professional activities

are macro design, micro operations, pattern recognition, and scenario formulation. We will see how and why the mess and reliability professional operates in the domain bound by the latter two hubs.

At the extreme of both scope and knowledge, where the scope is the whole system and the knowledge is formal, is the hub of macro design. Design—be it 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. Here design is meant to cover the operation of an entire system, including every single case relevant to providing system services. At the other extreme of the cognitive space is reactive behavior in the face of real-time challenges at the hub called micro operations. Here reliability depends on the immediate activities of system operators using tacit knowledge to manage a particular event rather than relying on preexisting designs at the system level for any eventuality. The field activities of crisis managers and emergency responders are micro operations par excellence.

Designers, however, cannot anticipate every eventuality. Worse, the more “complete” a logic of design principles aspires to be, the more

likely it is that its full set contains two or more principles that contradict each other—for example, we must not commit genocide . . . except when authorized to do so by the nuclear doctrine of mutually assured destruction.⁴ On the other side, operator reactions are likely to give the operator too specific or partial a picture, causing him or her to lose sight of the forest for the burning trees in the foreground. Micro operations, in other words, instill in us a kind of trained incapacity that undermines reliability because operators are not aware of the wider context(s) 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 has found that attempts to impose systemwide 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, if not illusionary. From the other side, an individual's reactive operations scarcely make for a widely tested template that can be applied to the system as a whole.

Instead of corner-to-corner movements, figure 3 indicates that reliability is enhanced when multiple shifts in scope are accompanied by multiple shifts in knowledge. Becoming more reliable means becoming more knowledgeable about varied things at variable scales. To that end, professionals approach and reach reliability through different skills than those for macro design and micro operations. Their approach is not direct, but indirect. System managers have to tack to reach reliability, much in the way that a sailboat does not get from A to B in a straight line, but rather frequently must cross into the wind to get there faster. To do that, however, requires knowing more than a straight line.

Specifically, 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 (the “scenario formulation and localized contingency scenarios” hub in figure 3). Food policy, for instance, works better when differentiating management protocols by crop or location (see Godfray et al. 2010, 813).

We also know that reliability is enhanced when operations shift from real-time reactions to recognizing patterns and anticipating their consequences across a run of cases of micro behavior and experience (the

“pattern recognition and anticipations” hub in figure 3). 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. Some patterns may be visible at one scale rather than others (see Schelling [1978] on micro motives aggregating into macro behavior). By recognizing and anticipating patterns across cases, operators and managers learn to adapt, and better practices emerge. 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. Signal detection and the ability to “read” feedback in terms of what these events indicate for the system as a whole are crucial for reliability management when operators don’t have full, immediate causal knowledge of the system they are managing.⁵

It is in this middle ground, bridging the formulation of design-inflected contingency scenarios realized more locally and the recognition of patterns and associated anticipations systemwide, that we find the reliability-managing professional networked with similar 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 system patterns then pertaining. In the middle is where the skills in pattern recognition and scenario formulation reinforce each other, as when repeated pattern recognition helps increase sensitivity to context-rich differences—and vice versa.⁶ In the middle is where reliability managers exercise their skills of interpretation as they translate pattern and scenario into managing reliably. And in the middle is where the reliability manager must be the mess manager.

For the middle is where we see operators and managers maneuver across the performance modes of the preceding chapter to ensure the safe and continuous provision of a critical service. This happens, moreover, with respect to services for which there are no formal control rooms or dispatch centers or trading floors, only networks of professionals to ensure that a service is provided reliably.⁷ This middle is, in brief, the domain of competence for these professionals. To say that managers are operating competently and skillfully is to say that they are managing within their unique domain of pattern recognition, scenario formulation, and interpretation so as to maneuver across the performance modes as conditions change. Otherwise, they could not be reliable mess managers; otherwise, they would not be the professionals that they are in terms of managing mess and reliability together. As

figure 3 indicates, no one else operating in the cognitive space of mess and reliability management has this unique knowledge base.

All of this sounds mushy, so some examples are useful. Within that space, the worst mess for professionals to be in—as first sketched in chapter 2—is to be pushed outside their domain of competence. You, the mess manager, are being asked to operate beyond your skills and talents. To be pushed from the credit derivatives you know to the derivatives of derivatives no one comprehends is a very bad mess. In contrast, the best mess is being well within the domain of known patterns and scenarios and translation skills, where managers can use their unique knowledge to maneuver across multiple performance modes under changing conditions. In other words, for those who take reliability seriously, good mess management is what occurs well within their domain of competence, with known but different patterns and known but different scenarios—sometimes with time to spare, sometimes with hardly a moment left. For these professionals, the chances of managing a mess badly increases the closer it is to the limits of known patterns and scenarios—that is, the closer it is to the edge of the domain of competence (or what Paul Schulman calls the precursor zone). It is here where reliable mess managers are short of knowable options and where the options they have could have important but unknown effects.

This management is, again, without any guarantees. Bridging scenarios and patterns, each of which differ with the others, is the difficult part of the professionals' translation, because the interpolation involves transposing, transforming, and synthesizing scenario and pattern in light of others in order to manage in the present.⁸ Translation is needed if only because localized scenarios and systemwide patterns are drawn from very different knowledge bases (figure 3). Again, that translation is interpretative rather than literal, and this is how new or different knowledge is generated—though that process is not without its own risks, as we have seen.

But just what exactly is involved in “translation” or “synthesis” of patterns and scenarios? One way to start thinking about this is to recognize that the macro designers and micro operators around the two extreme hubs who talk about “risk,” “coordination,” and “learning” are frequently doing so differently from those in the middle domain. If the differences were appreciated by decisionmakers, those endless debates over planning versus implementation or comprehensive planning versus piecemeal incrementalism, among others, would have to be rethought. Let's introduce each briefly here, leaving fuller comments to later.

When professionals in the middle talk about objective risks in large technical systems, they mean not only worst-case scenarios that they have formulated or the hazards and frequencies that they have recognized. They also mean the risks that emerge out of the unique knowledge it takes to manage the systems across and within the four performance modes—each of which has its own dominant risk. Activities in the middle domain are not more certain and less risky (or less messy and more reliable); rather, the operating complexities and risks are multiple and change with shifts in scope of management and knowledge needed to manage reliably in that domain.⁹ The professionalism comes in knowing these differences and why reliability is risky in the ways it is.

Learning is also different for those looking from the middle to outside the domain of competence. The domain in figure 3 is not static. Patterns and scenarios are added to or dropped from the repertoire, as messes and the professionals who manage them adapt to changing circumstance. By implication, not only can macro-design “solutions” that bypass this learning pull professionals and their networks outside their middle domain of competence, but the interventions just as often fail to capitalize on the evolutionary advantage of these middle mess managers in improving reliability operations and rejiggering processes and technology to ensure those improvements.

“Coordination” must also be rethought. By calling for greater teamwork or stakeholder coordination, mess and reliability professionals in the middle often mean “bringing the system into the room” (Weisbord and Janoff 1995). To do so is to bring in those with expertise in macro design and micro operations, but also those who network patchy patterns and scenarios into reliability. This ability to reconnect disconnected activities in ways that better match or mimic the connectedness of reality is at the heart of the professionals’ translation.

If you look closely at figure 3, you will see that we are talking about professionals who are experts not because they “bridge” macro design and micro operations directly. On the contrary, professionals synthesize knowledge about planning and about operations into reliable services, however messy their translation may be in the face of uncertain success. One great mistake in conventional policy analysis and public management has been to assume that implementation is all about converting macro design into micro operations or that implementation at the micro level ends up as a kind of *de facto* policymaking at the macro level. Nothing could be further from the truth. Implementation takes place in the middle across a network of professionals. 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 policy¹⁰—to networks of middle professionals and the risks they face. It is there where messy trade-offs and interpretative muddles are to be found between the pattern recognition and scenario formulation hubs, and where any better practices that emerge across a run of micro operations have to be modified in light of local contingencies.

Stay with those street-level workers for the present, because the differences between them and the middle domain of mess managers help us to understand just what the latter professionals actually do by way of managing for service reliability. Each of the two groups is oriented differently to the hubs and domain of competence, and the differences in orientation are instructive—although nothing is hard and fast here—when it comes to understanding the nature of mess management as discussed in this book.

First, there are differences with respect to pattern recognition. At best, street-level workers avoid labeling and stigmatizing clients: “Street-level workers do not see citizen-clients as abstractions—‘the disabled,’ ‘the poor,’ ‘the criminal’—but as individuals with flaws and strengths who rarely fit within the one-size-fits-all approach of policies and laws” (Maynard-Moody and Musheno 2003, 94). But in my extended framework, pattern recognition differs considerably from a macro design of one size fits all. For the mess and reliability professional, stereotyping is its own systemwide phenomenon, with its own patterns. For example, how do people vary in terms of education, age, ethnicity, income, or gender when it comes to stereotyping? Mess and reliability professionals want to know the better practices for dealing with such stereotypes when it comes to “juveniles” or the “disabled.” Street-level workers have to first know persons, though they too stereotype from time to time; mess and reliability professionals have to first know how populations differ, though they too work one-on-one from time to time.

There are differences with respect to one’s stand toward macro design. For the mess and reliability professional in a network, macro design is as disputed, incomplete, uncertain, and complex as any other hub in his or her operating space (we will see this in the overpopulation example that follows). There is no one overarching morality or standard when it comes to trying to avoid trade-offs. That is why professionals connect with other professionals in order to get anything done halfway reliably. For the street-level worker, the moral order is clearer: “For example, the decision to subvert the rules by an exasperated [social

service] counselor . . . redeems the state by breaking through the bureaucratic labyrinth” (Maynard-Moody and Musheno 2003, 24). Street-level workers may be willing to subvert departmental protocols and procedures in order to do the right thing. Mess managers in the middle managing for reliable critical services face a greater variety of plural values when it comes to rights or wrongs.

There are differences with respect to where the street-level worker and the middle professional stand with respect to localized scenarios. The street-level worker may have a more negative view of localized scenarios than the mess and reliability professional does, when those scenarios are protocols and rules devolving from departmental policy. For the street-level worker, macro policy and localized rules are much the same thing: the problem and a cause of difficulty. For the reliable mess manager, those localized rules are resources to be exploited in order to keep departmental services reliable, as task conditions change.

Differences in orientation to micro operations are also notable. For mess and reliability professionals, the individual case is a starting point from which to search out patterns over a run of such cases. How else do you find better practices? For the street-level worker, the individual constitutes the center of gravity of service provision. Numbers, trends, and procedures are really not the endpoint; the worker’s relationship with the client is. “Indeed, the worker’s decision of when to conform to rules and procedures and when to break them and when to cooperate with authority and when to act independently is the essence of street-level judgment” (Maynard-Moody and Musheno 2003, 68).

Finally, there are also differences in the stand the two groups take with respect to what I have been calling “the middle.” For street-level workers, the middle drives the system and is very much part of the problem: “In their stories, the system is described as an undifferentiated amalgam of other units in their agency, other agencies, elected officials and the media. . . . Street-level workers see themselves as moral actors working in opposition to the system and rarely describe themselves as part of it” (Maynard-Moody and Musheno 2003, 22). For the mess and reliability professional, the middle is far more differentiated, set as it is between the hubs that bookend it—namely, localized scenarios and recognized systemwide patterns that rely on different mixes of knowledge.

To summarize, for the reliable mess manager, patterns and better practices matter as much as protocols and procedures, and it is within networks that these are to be managed. For the street-level worker, face-to-face relationships matter more than protocols, and headquar-

ter networks are power elites to be circumscribed, when not circumvented. When the latter happens, the street-level worker can be part of the bad mess in which middle mess managers find themselves. From the other side, when there is no network of reliable mess professionals in the middle (assume that they're all operating in unstudied conditions), the street-level worker is indeed alone, acting in ways that necessarily equate professionalism with reliable micro operations.

Before next turning to a specific policy application of the framework, a preceding point must be highlighted. Where you see one mess and reliability professional, you see a network of them. Policy messes are so complex that a reliable mess manager in the middle cannot manage any one of them on his or her own and still be reliable. There has always been something dangerously misleading in public policy and management literatures that perpetuate deracinated notions of "policy entrepreneur" or "change agent," as if each were the counterpart to the solitary street-level worker.

An Application

Arguably, the world's most important mess when it comes to public policy and management has been the long-standing controversy over global overpopulation and associated overcrowding.¹¹ The crisis narrative is a familiar one. Human population numbers—some seven billion people with a net increase of over seventy-five million a year (see, for example, Bloom 2011; Wolf 2003)—threaten our planet with unprecedented overcrowding, environmental spoliation, and resource conflicts. We are fast approaching, if we are not already past, the sustainable limits of water, clean air, and energy. Without population restrictions, including but not limited to birth control and growth limits on cities and all manner of resource utilization, the globe is headed for irreversible decline, assuming that has not already been assured. *The Long Emergency* (Kunstler 2005), not *The Long Boom* (Schwartz, Leyden, and Hyatt 1999), is under way.

Consider the numbers, we are urged. In the early 1950s, global population was predicted to be 3.6 billion by 2000; the actual figure was more like 6.1 billion (R. Cooper and Layard 2002, 8). The planet's population has been forecasted to reach up to 12 billion by 2050 (9). What about our natural resources? The Food and Agriculture Organization of the United Nations calculates that forest loss has been huge: a net loss of 6.4 million hectares between 1990 and 1997 alone (Kaiser 2002, 919). Projections for

energy and water use signal terrifying depletions (Brown 2002). More than half of the world's population turned urban in the first decade of the twenty-first century; a century or so earlier, urban populations represented less than 15 percent of the total (see, for example, Crossette 2002; Greenhalgh 2010). Surface temperatures have risen over the last century, and global climate change continues unabated (Ramanathan and Barnett 2003). The world's greatest problem is population growth, according to James Watson and Francis Crick, the discoverers of DNA (Daugherty and Allendorf 2002, 284). What, Jared Diamond asks in *Collapse* (2005), was that Easter Islander thinking when cutting down the island's last tree? Two conservationists in *Science* conclude: "One word sums up the overall and long-term problem [in creating a sustainable future]: overpopulation. We wonder how any sane person could disagree" (Wright and Okey 2004, 1903).

Here is how sane people disagree. First, the numbers are disputed, and population projections remain full of uncertainties (see, for example, Walker 2009). The United Nations revised its global population projections substantially downward at one point, and it was estimated that the total would be around nine billion by 2050 (see, for example, Chamie 2010, 157; United Nations 2003). The figures were subsequently revised upward, to just over ten billion by 2100 (Gillis and Dugger 2011; R. Lee 2011). "There is, however, considerable uncertainty surrounding these projections," as a professor of economics and demography insists (Bloom 2011, 562). One study indicates a forest loss of 20 percent less than the original Food and Agriculture Organization estimates, while major water-use projections have been overestimated (Brown 2002; Kaiser 2002). The only certain thing about global energy projections is that they are wrong, if we believe the experts. Temperatures have been increasing, but a vigorous controversy continues over what this actually means regionally and in terms of costs and benefits (again, start with Ramanathan and Barrett 2003). Instead of focusing on that Easter Islander, we might just as well ask what European explorers thought they were doing when they knowingly introduced venereal diseases to the Pacific (N. Thomas 2003, xxv–xxvi). Finally, and with all due respect to Watson and Crick, when did they become experts on population growth?

Other problems with the data and methods must be registered. Strong taxonomic biases in documenting species have long existed in conservation research (Clark and May 2002), estimates of biodiversity losses remain disputed, and there are those who consider global urbanization to have net benefits for controlling total population numbers—for example, family size and birthrates tend to drop when populations

become more urban (Revkin 2002). We must, of course, tread cautiously here. I am not saying that there is nothing to worry about by way of overpopulation—no one wants some ecologists equating him or her to a Holocaust denier.¹² Does this mean, then, that analysts are wedged between two conflicting narratives with respect to population and crowding, waiting for the evidence to free them?

Figure 3 suggests how to sort out this policy mess. Explicit in the call for population curbs, particularly those limits on growth and births, is the macro-design concept of a global carrying capacity, which is the upper limit or cap on the total number of people that the planet can support without collapse. The idea that the level of sustainable population can be derived from a calculation of global carrying capacity is contentious on several counts, however. First, which global carrying capacity estimate do we rely on? At the time of writing, the only certain number has been sixty-nine, which is the number of past studies reviewed in a meta-analysis of the widely divergent estimates of global carrying capacity. The meta-analysis found the lower and upper population bounds were 0.65 billion and 98 billion people, with its best-point estimate of 7.7 billion (van den Bergh and Rietveld 2004). Second, major ecologists doubt whether there is a “carrying capacity” for arid and semi-arid lands, which constitute much of the surface area of the planet (Roe 1999; see also Scoones 1996).

As for the other extreme, the micro operations of overcrowding—the actual experience of overcrowding—are full of distinctions. What feels overcrowded to someone in Europe need not be so to someone in Southeast Asia. What feels overcrowded to rural residents may not to urban residents within the same country. Even when both sets of residents concede that their areas are overcrowded in the same way, one group might say the solution is not fewer people as much as it is more education or technology. Even if they agreed that their areas were overcrowded for exactly the same reasons with exactly the same effects, it is unrealistic to believe that anyone knows enough, no matter what his or her expertise, to recommend what the actual population levels should be for a wide area concerned. It is difficult enough for a long-term resident to make such judgments for his or her smaller locality, let alone the most complex ecosystem there is: the planet.

As we tack from macro design and individual experience in figure 3, we add to the knowledge bases about population, age structures, population densities, and related factors. We have already seen that the very different global trends and generalizations do not match the dominant, more uniform macro narrative about global overpopulation

(that is, we have noted the differences between carrying capacity as a macro-design concept and the actual variability in estimates across the globe at the pattern-recognition hub).

Substantial differences also emerge when we move the analysis from macro design to localized contingency scenarios. Stay with global carrying capacity as core to the determination of overpopulation at the macro level. When we move from that governing concept to its regional counterpart, it turns out that a handful of the world's regions—most notably India, China, Pakistan, Bangladesh, and Nigeria—have accounted for recent major increases in world population (Wolf 2003), and other substantial regional differences persist (Roberts 2011). Indeed, birthrates have started to decline in China, while fertility rates in Bangladesh are projected to decline (Walker 2009). So what are the carrying capacities of these specific countries, and who knows enough to give that answer for the next twenty or more years?

Once we stay focused on conventional regions of the world, Europe comes readily into view, as a big policy mess there has been shrinking population levels and declining fertility relative to health and social service demands (Lutz, O'Neill, and Scherbov 2003; Ringen 2003; see Walker 2009 for intra-European differences). Other regionalized protocols and scenarios come to the fore as well. Concern for climate change at the global level has moved to developing and improving regional climate models, just as national weather models have become more regionalized (Kerr 2004). There are clear regional differences in climate changes and their effects on species, for example (see Myers and Pimm 2003).¹³

Now, let's plot these positions and findings for the overpopulation and overcrowding controversy in a mess and reliability space (see figure 4). The dimensions and plot of positions in figure 4 help us to answer these questions: Just who are the mess and reliability professionals in this controversy? Who is competent enough to move across the four hubs and translate the system patterns and regional scenarios into reliable service provision (be they for water, air, or energy) that are said to be challenged by population growth and overcrowding?

Whoever the middle professionals are—again, they most certainly are not all in control rooms—they have different knowledge bases than those at the extremes of macro design and micro operations. They must work somewhere between the regional (localized contingency scenarios) and the global (system pattern recognition). We must also expect that the professionals and their networks are already there—they do not have to be created from scratch. For highly controversial

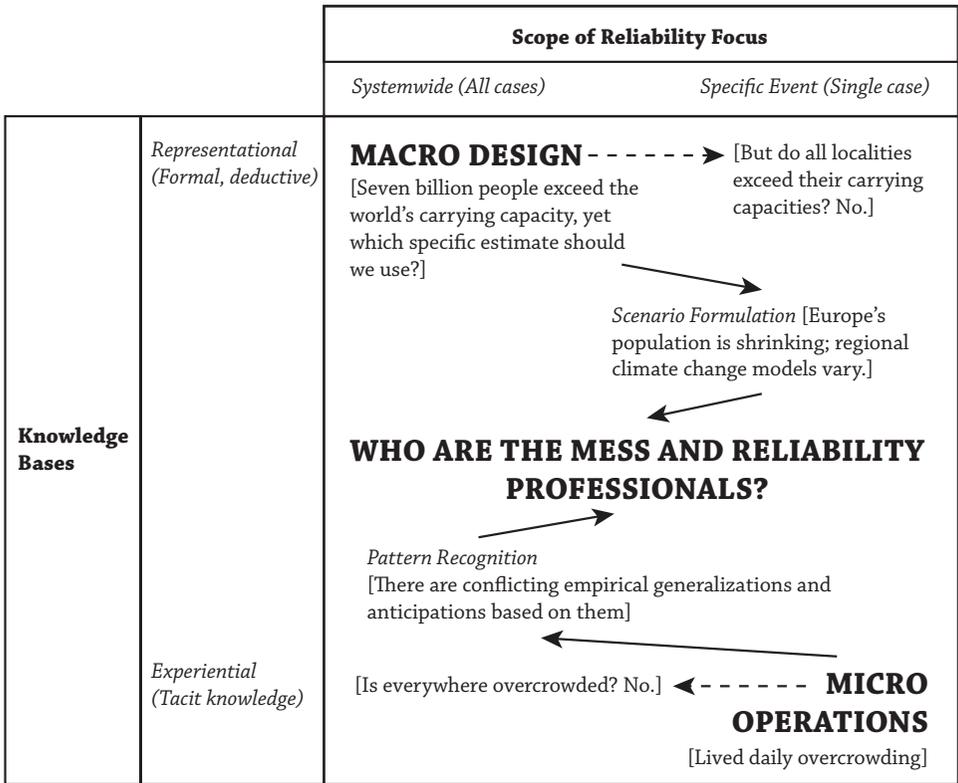


Figure 4. Mess and reliability space for the policy mess relating to global overpopulation and overcrowding

Source: Adapted from Roe 2007

policy messes like the one regarding overpopulation and overcrowding, we can assume that almost all points in the mess and reliability space are occupied.

But why are we interested in these professionals? It is not because they have the “solution” to overpopulation or overcrowding, but because the policy messes around the four hubs have to be translated, if possible, into unique knowledge for securing more reliable services in the face of all manner of population and crowding pressures. Look again at figure 4 and move to the middle from the localized-scenario and pattern-recognition hubs. What falls between a major region, formally “Europe” or “Southeast Asia,” and the globe? One familiar answer is the classic nation-state. Choose two countries whose populations are treated as having similar systemwide patterns and localized scenarios when it comes to this policy mess—that is, the citizens of

both countries by and large say they are overpopulated and overcrowded and will become increasingly so. In this way, the countries are taken to share the features purported to exist in the dominant crisis narrative about worldwide overpopulation and overcrowding. For example, choose the Netherlands and Singapore. Ask people in both places “Is this country overcrowded?” and you would probably get a qualified “Yes.” The population densities are perceived to be very similar between the highly urbanized western Netherlands—the Randstad—and Singapore (Roe and van Eeten 2001). Now the question to ask is: Based on experience and familiarity with the two countries, what policies enable their residents to accommodate their human populations? Or bluntly, when the countries are perceived to be overpopulated, why aren’t more and more of their citizens being pushed outside, beyond the national borders?

From a management perspective, a country is overcrowded and overpopulated when its mess and reliability professionals have few ideas about how to keep people residing, employed, and productive there. Countries can move in and out of conditions of over- and underpopulation and over- and undercrowding, depending on the livelihood strategies adopted by their residents. This means a country that is not overcrowded can become so, even if population numbers or densities do not change. All that needs to happen is that the management of its policies worsens. This suggests that an explicit management goal of economic and social policy of a country or region should be one of retaining and sustaining people who are already there and want to stay. The Netherlands and Singapore risk becoming (more) overpopulated and overcrowded only when increasing numbers of residents there want greater well-being yet choose to leave, even if they are uncertain as to whether their greater well-being lies somewhere else. They are pushed out, rather than pulled elsewhere.

This translation of the policy mess of overpopulation and overcrowding stands in sharp contrast to current orthodoxy about global population increase. Our new policy narrative does not claim that overpopulation and overcrowding are not problems; then again, it does not claim that one country’s medley of policies for addressing population and crowding will or should work elsewhere. There is a story here, but it is unique and does not have the same ending for everyone. It is a contingent narrative, provisional on how the networks of professionals translate the patterns and scenarios involved for where they are.

What the new policy narrative does claim is that the dominant narrative and those who criticize it without offering alternatives avoid, dis-

regard, or dismiss the middle domain of competence. The dominant crisis narrative insists: System trends are this way, and therefore macro solutions must be that way. But there is no “therefore” when we are in the middle, looking out. Why ever would we jump from pattern recognition to macro design without first consulting those who are already addressing population and crowding in varied and instructive ways? We return to and expand on this point in the next two chapters, which discuss what makes for bad and good mess management.

It should go without saying that the mess and reliability professionals who deploy this translation of an overpopulation and overcrowding narrative vary from country to country and cannot be equated to any single cadre, such as macroeconomic planners. The professionals I’m talking about are the ones who excel at cross-scale, context-dependent, case-by-case analysis. When it comes to dealing with population and crowding, they are the ones who are adept at finding several ways of achieving the desired policy ends, albeit sometimes with little time to spare and without “proper” planning. They bring others into meetings so that the entire system is better represented. They seek to recouple what has been decoupled through disciplinary specialization, program fragmentation, and agency turf battles. They continue to search for better practices to avoid losing options; they distrust estimates of hazard when it is not based on intimate knowledge of the system to be managed; and they frequently strive to achieve flexibility—though they never get as much slack as they can reliably use—in the provision of services that society considers critical. Of course this is messy, but there is no other way to be as reliable.

This recoupling through management of what has been decoupled organizationally is critical to professional mess managers. But how does it work? I once attended a presentation on an ecosystem restoration project in Montana. A leader of the project described his take on its key contribution: He could now see how his forested acres fit into the wider valley landscape. The approach gave him a way to integrate the small and large scales, with cross-scale implications both ways. Not only could the project leader stand in his woodlot and see its role within the larger scale of the ecosystem and landscape, but he was able also to plan at the smaller scale for the longer term. Similarly, as the ecosystem manager stands at the ridge looking down into the valley, she is able to plan at the larger scale for the shorter term. She can now see what the next steps are when it comes to managing the entire ecosystem. “Think globally, act locally” becomes “Think long term from the small scale, act real time from the large scale.”¹⁴ If and when

this happens (again, no ironclad assurances here), recoupling—as messy as it might be—can end up reworking policy and management.

The difficulties with recoupling lie in searching far and wide for better practices and in having both the requisite contextual knowledge so as to formulate a local scenario and the know-how to modify the better practices in light of that local context. In this way, “context” is brought into the analysis not just at the aggregation stage across a run of cases (what we have been calling pattern recognition) but also in the modification of these practices in light of the relevant local scenario(s). In terms of figure 3, we are in the domain of professionals, seeking to apply and modify systemwide better practices with respect to ecosystem management to the local scenario at hand, which is rich in its own specifics. Professionals ask, “Who has figured it out better, and how can that be modified for use here?” This is how small actions add up to big effects: Local actions are based on 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. The great advantage of learning this way is that better practices import the scales at which management actually works across a run of cases without having to prejudge what are “the right planning areas.” If this sounds like old-fashioned incrementalism—though it is nothing like the version I was taught—then call it “Incrementalism,” where the capital *I* indicates a scope of search for better practices that aspires to be truly international.

It should go without saying that this process of feedback and updating better practices is dynamic: There is never a “best practice,” and even the economists’ default to second-best solutions may still be far too simplistic. Current better practice is not some kind of macro standard to which we all are meant to aspire. This is not about measuring any one-meter stick against the hermetically sealed international standard meter bar at the Bureau of Weights and Standards outside Paris. Better practices are more akin to pointing out how this heart I am looking at on the screen compares to those digitized ones that I am using by way of comparison. Where does this heart fit into those that are said to run the gamut from healthy hearts to unhealthy ones? Too many decisionmakers look for the equivalent of platinum bars in policy and management while ignoring the instructive specimens in front of them, each a messy original on its own but sharing family resemblances with others. In the language of narrative analysis, better practices transform patterns into stories that can be used by decision-

makers. To put it differently, better practices ensure that the emergent behavior of micro operators is not reducible to “just statistics” but has meaning for management in ways that substantively differ from those stories told by macro designers, including policymakers and lawgivers.

Speaking of which, what family resemblances do this framework and its implications identify for conditions leading up to and including the 2008 financial mess?

Turning Back to the Financial Mess

It takes only a few small steps in the crisis narrative to move from insisting that there are too many people on the globe to concluding that these people are now globally interconnected, and importantly so. In this logic, more people mean more resource scarcity, and fewer resources mean more interdependence, and more interdependence is just what has been happening through globalization policies that integrate economies together—and their financial systems. Many such examples exist: At one point, for instance, the announcement of even more quantitative easing in the United States was followed a few hours later by the Mumbai stock market moving to its highest level in nearly two years (Sri-Kumar 2011). A *Financial Times* correspondent deduces: “There is a strong case to be made that the current [financial] crisis is in the strictest sense a crisis of globalization” (Guha 2009a).

What does our framework have to say about this line of argumentation? Start with the major macro-design position on financial globalization prior to the events of late 2008: The integration of a country’s financial markets into the world’s economy increases that country’s economic growth and worldwide economic growth. Given figure 3, we can expect that other positions in the mess and reliability space for globalization would differ substantially from this macro position, and that was the case well before the financial upheaval. First, economic growth under always-late capitalism has never been uniform across the globe. There really is no one “protocol” for economic growth worldwide. According to the head of the Global Economic Research unit at Goldman Sachs, since 2001 (but before the financial mess), the handful of BRIC economies—Brazil, Russia, India, and China—“contributed about one-third of all global growth” (O’Neill 2006, 12). Not only does the macro-design position for financial globalization not conform to the existing but differing scenarios for emerging economies, but the systemwide patterns and trends also vary significantly from the macro

standpoint. Prasad and Rogoff reported earlier on the findings of a study on globalization's effects:

Interestingly, the more financially integrated developing countries do seem to have achieved higher per capita incomes than others. However, it becomes difficult to make a convincing connection between financial integration and economic growth once other factors, such as trade flows are taken into account. . . . We found that financially integrated developing economies have in some respects been subject to greater instability than other developing countries. (2003)

As for examples among developed countries, there never was one “European social model” from which to design individual welfare systems. Rather, a handful of social models were discernible in the empirical evidence across Europe, none of which applied directly to any single country's scenario without a good deal of translation (see, for example, Sapir 2005). Certainly, when it comes to the global financial mess, some regions and countries—notably Canada and Australia, as we shall see—came out relatively unscathed.

Let us stay with the difference between the pattern-recognition and macro-design hubs a moment longer, because even thoughtful commentators on globalization conflate the two. Martin Wolf took this to be one of his “ten commandments of globalization”: “It is in the long-term interest of countries to integrate into global financial markets. But they need to understand the need for an appropriate exchange rate regime, often a floating rate, and a sound and well-regulated financial system” (2004). This statement, cast as a principle, is at best an empirical generalization contingent on a reading of the weight of evidence—that is, pattern recognition—at the time. Patterns, however, are not macro principles, and we will see in chapter 4 that such a conflation of hubs can lead to very bad mess management. While it would be a fairly simple matter to continue showing how positions in the globalization space differ, the core issue circles back to these questions: Who are the mess and reliability professionals in the middle who can make sense of this globalization space? Whose task is it to reconcile these conflicting and complex positions in the name of reliably managing globalized services, such as finance?¹⁵

Let's return to the media reports about the 2008 financial upheaval. You would think the real managers in the global financial mess were senior officials in Treasury departments, central banks, and finance ministries. In our framework, those decisionmakers are better understood to have been macro designers acting reactively as micro opera-

tors. If so, this should be troubling to anyone interested in the reliability management of major financial institutions. We cannot and should not expect senior officials to be the primary professionals who actually determine whether the bailout and other interventions achieve and sustain reliable real-time financial services, as messy as the aftermath of the financial meltdown must be. Baron de Montesquieu, the French political thinker who was also criticized from many sides, once compared himself to someone living on a second floor “disturbed by the noise upstairs and the smoke downstairs” (quoted in Todorov 2009, 27). Caught as they are between loud macro designers and firefighting micro operators, many reliability professionals in the financial mess doubtless felt and still feel the same.

Who, then, are the real-time managers and operators called mess and reliability professionals when it comes to the financial upheaval? We introduced them in chapter 1, but the framework is now in place that connects them and underscores the seriousness of their role. They are the members of that unique class of professionals—especially middle-level managers and support staff—whose supervision, informal networks, and skills ensure that financial services do not fail as often as they could. As we have discussed, you find them in IT units, accounting units, engineering divisions, line operations, business continuity staff, inspectorates and supervisory units, auditing departments, and regulatory and legislative offices as well as on trading floors and in the field, and not just working in the area of financial services.

Some in the executive office suite, like the chief financial officer or immediate staff, might be part of the network, though they are scarcely leading all the real-time operational decisions involved in managing reliably. Occasionally the middle-level staff and specialists appear in the press as self-identified “plumbers” (Grant 2009b; O’Connor 2008), but they rarely surface to the public’s attention—and when they do, even more rarely are the networks in which they work made visible. In fact, some call this managerial and networked know-how “dark matter” (Hausmann and Sturzenegger 2005). This means that there is absolutely no reason to believe all major executives have this operational knowledge, though some of the founders of the United States, for instance, insisted that it should be otherwise.¹⁶ During the financial mess, we witnessed the tribulations of a private equity investor in a major car firm who, by one account, “did not have a clue about the automobile industry” (Story 2009); the partial dismantling of another major automobile firm by “a not-quite graduate of Yale Law School who had never

set foot in an automobile assembly plant” (Sanger 2009); and the new chairman of that firm who admitted, “I don’t know anything about cars” (Gapper 2009). Also in the period leading up to the financial mess, we witnessed the problematic role that consultants played when they saw their task as one of macro-designing system changes rather than supporting networked professionals already in those operations. For example, Citigroup’s “reliance on outside consultants for strategic advice on credit instruments contrasts with the practices of rivals that fared better than Citi in the crisis, such as Goldman Sachs and JP Morgan”; the latter relied on in-house talent before and during the meltdown (Guerrera and Politi 2010). For the time being, think of this kind of know-how as one depending “on complex mixtures of judgment, problem-solving and information exchanges, often involving group behaviour that is difficult to replicate” (Manyika 2006, 13).¹⁷

Summing Up

We are now positioned to say much more than that a policy mess is an amalgam of uncertainty, complexity, conflict, and unfinished business, or that it varies in terms of the performance conditions. This chapter’s framework enables us to define and summarize more formally what a policy mess is and the management it entails.

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. For our purposes, a policy controversy is not really a mess until different and conflicting positions across the different hubs are taken on the issue. Management is to sort out the different positions at each hub and across hubs, and the following two chapters show how that can be done poorly or done well. Whether reliable mess managers in networks of like professionals can extract a good policy mess from a bad one, or stop a policy mess from going bad, depends on their unique knowledge of how to synthesize patterns and scenarios into reliable service provision. Pattern recognition and scenario formulation are, if you will, the oxygen of that management.

In a policy mess, gaps in knowledge always exist between macro design and pattern recognition, between scenario formulation and micro operations, and between pattern recognition and localized scenarios. A gap does not mean the respective hubs are polar opposites, only that different blends of deductive and inductive knowledge sepa-

rate them. When it comes to mess and reliability management, what a manager holds at the level of macro theory for the system and what that professional finds in practice at the system level *necessarily* differ. Similarly, system patterns and the anticipations based on them inescapably differ from local scenarios that seek to contextualize design considerations better. It should never surprise any serious manager to find a gap between what a regulation says formally and how its requirements are realized in a given region or between that formal regulation and what is actually found in emerging practice across multiple regions. In fact, to try to ignore such gaps is to turn reliability management into bad mess management, thereby undermining the critical service these professionals are trying to provide.

Finally, I have been writing as if policy messes are good, bad, or ready to go one way or the other, depending on how they are managed. But, as we already saw, some messes are good and bad at the same time, and that fact raises a question we should be asking about all policy messes: Just what are we managing for? The painter Gérard Fromanger points out that a blank canvas is white but also “black with everything every painter has painted before me” (quoted in Shatz 2010)—and so too there are messes that are both bad and good in the same instant. Bad policy mess: It is said that one out of every two young African-American men in major urban areas is enmeshed in the criminal justice system. Good policy mess: Why, then, are we not interviewing the other 50 percent of young urban African-American males outside the criminal justice system to find out what they are doing, and what the rest of us could learn from them?¹⁸ Bad policy mess: At one point, three to four billion people—up to two-thirds of the world’s population—lived in regions without adequate water supplies or sanitation (see, for example, World Health Organization 2007), a dire situation for development agencies and experts to address. Good policy mess: Now that is a very large number of people, right? This is such a huge distribution of people without adequate water supplies, that some of them must be doing much better than the others. That means then there are tens of millions—hundreds of millions?—of people who have many things to say about how to better survive without adequate water to those millions more who are also trying to survive without it.

So the next time someone complains, “The economy is in a mess,” press them a bit: Is it that trends observed at the systemwide level are at historic lows; or that the trends differ from some design optimum, as when economists talk about the output gap; or that individuals they know are having a harder time? Or is it that, while all this is going on, a

good part of the mess lies in the hodgepodge of bustling economies mixed in with deep-recession others, depending moreover on how “the economy” is defined locally?

That good and bad mess can go together—and that the task may be to pull the good mess out of the bad—is also illustrated through the example of one favorite whipping boy of public policy: government subsidies and interventions that encourage redevelopment of areas that have repeatedly flooded. A flood breaches the levee; the houses, roads, telecommunications, sewers, and electricity lines are destroyed—and what does the government do? It rebuilds the levee so that the redevelopment of critical services starts all over again! Really, how dumb can the government be, right?

Wrong. For my part, I would also want to know if the flood enabled the telecommunications provider, for example, to replace legacy equipment it could not replace otherwise due to prior regulatory and insurance considerations. I’d want to know what eventuality that provider prefers: (1) taking advantage of the opportunity to replace out-of-date or unprofitable equipment it could not replace except as a result of emergency action during the “bad times” of a flood; or (2) having to maintain its market share by severe cost-cutting for equipment maintenance, repair, and replacement, as a result of man-eat-dog competition during the “good times” between floods. Such two-sided policy messes are extremely important because their pushes and pulls force decisionmakers to focus on asking what we all should be asking of any policy mess, be it in health, environment, social services, finance and banking, or another arena: Just what mess are we managing, and how are we undertaking that management? The next two chapters answer those questions.