11

DATA INDETERMINACY: ONE NASA, TWO MODES

Roger Dunbar and Raghu Garud

In the early morning of February 1, 2003, the incoming *Columbia* shuttle flight 107 started re-entry into the earth's atmosphere. At 8.54:24, however, Mission Control noticed that the hydraulic sensors on the left wing had failed. The spacecraft was committed to re-enter the earth's atmosphere, was traveling at around Mach 23, and its wing temperature was set to rise to over 2,800 degrees Fahrenheit. Mission Control simply watched as the spacecraft disintegrated.

An investigative board highlighted the importance of a specific technical event during the shuttle launch, i.e., a block of insulation foam around 640 cubic inches in size fell off and struck the underside of the orbiter's left wing, breaking several protective tiles and most likely compromising its thermal coating skin. When the spacecraft re-entered the earth's atmosphere 16 days later, hot plasma gas flowed directly into the spacecraft, setting off chain reactions that destroyed the shuttle's systems and crew. Several NASA managers had realized during the launch that this destruction might be possible. They asked that some checks be carried out, but most were never done.

Although the Accident Investigation Board identified the specific technical events that caused the orbiter's physical destruction, it also suggested that the tragedy involved critical organizational issues. Since the earliest shuttle launches, blocks of foam had fallen and struck the orbiter, potentially damaging its thermal protection system. Despite these events, NASA had continued to fly shuttles, including *Columbia*. What were NASA's managers doing?

Vaughan (1996) described how, over time, a process she calls "normalization of deviance" became a NASA routine. Deviance normalization accepts that unexpected events occur. It then requires NASA managers to classify unexpected events as being "in-family" or "out-of-family" and involving "not safety of flight" or "safety of flight" issues. By defining "in-family" and "not safety of flight" issues as "acceptable risks," managers can permit launches, helping NASA maintain its flight schedule commitments.

In this chapter, we deconstruct the processes and actions that impacted the ill-fated STS-107 shuttle (the *Columbia*) flight. Our general argument is that alternative organizing modes for managing distributed organizational knowledge ensure data indeterminacy. We begin by explaining how knowledge in organizations is distributed across different organizational elements, making alternative organizing modes possible. One organizing mode emphasizes the use of knowledge in support of exploration. A second organizing mode emphasizes the use of knowledge in support of predictable task performance. Depending on the organizing mode that an organization relies upon, pieces of knowledge are pulled together in different ways to generate an organizational response.

We use the story of shuttle flight STS-107 to show how, when people process real-time data, ongoing events have different meanings in different parts of an organization. We also highlight how understandings that are appropriate for achieving predictable task performance may be directly at odds with understandings that are appropriate for exploration. We show that as individuals operating in real time attempt to accommodate both organizing perspectives simultaneously within a distributed knowledge system, the significance of available real-time data becomes indeterminate so that ways to react or respond become impossible to discern. In situations that demand high reliability because of the potential high cost in human life, the emergence of indeterminacy can have disastrous consequences, as was the case with STS-107.

DISTRIBUTED KNOWLEDGE: PEOPLE, TECHNOLOGIES, RULES, AND METRICS

Where exactly is knowledge and understanding located within an organization? There is growing acceptance of the idea that, rather than being somehow shared in organizations, knowledge is fundamentally distributed across different organizational elements (Hutchins, 1995). It is distributed not just across and within people but also in organization technologies and their designs, in the rules and procedures that are used to identify contexts and mobilize actions, and in the metrics and tools that are used to determine and assess value (Garud and Rappa, 1994; Callon, 1998). Table 11.1 identifies a set of organization elements each of which contains different pieces of distributed organizational knowledge.

If an organization is to use its knowledge to generate value, the elements that contain distributed knowledge must somehow cohere. Coherence evolves as people in organizations choose specific metrics and install procedures to use and combine knowledge from different elements to achieve value according to these metrics.

Table 11.1 Distributed arrangements

People	Different perspectives and different levels of inclusion
Technologies	Knowledge is embedded in technological artifacts
Organizational routines	Establish the decision context and the temporal rhythm for coordination
Metrics	Shape what is measured, what is acceptable and what is not acceptable

Metrics and contribution combinations change, however, as new tasks appear, and new tasks often appear as procedures applied to technologies achieve unexpected results. Metrics may also change due to outside interventions, as when people representing new interests are brought into an organization. An organization's distributed knowledge is able to combine and cohere in an evolving way, enabling what an organization does and how it does it to adapt and change.

Given that metrics have been agreed upon, Tsoukas observed: "The key to achieving coordinated action does not so much depend on those 'higher up' collecting more and more knowledge, as on those 'lower down' finding more and more ways of getting connected and interrelating the knowledge each one has" (1996: 22). Effective coherence requires not just that people as knowledgeable agents cooperate but also that they and an organization's technologies are embedded within a larger set of procedures that reflect broader knowledge sources enabling and constraining action so that an organization achieves value according to its metrics. Effectiveness depends on how the knowledge distributed across all organizational elements is pulled together in a particular situation.

LINKING ORGANIZING MODES AND DISTRIBUTED KNOWLEDGE

An organizing mode is a normative orientation that facilitates a particular way of using organizational knowledge. An organizing mode highlights what is significant for an organization, mobilizes energy to facilitate what is significant, and enables sanctions to maintain what is significant. Specifically, organizing modes direct that knowledge located in an organization's distributed elements – people, technologies, procedures, and metrics – should be used for a particular purpose in a specific and mutually reinforcing way.

Organizing in exploration mode

To illustrate how an organizing mode influences the development and use of knowledge, we consider 3M Corporation. As innovation is the central aspect of 3M's corporate identity, 3M's organizing mode for managing distributed knowledge should consistently support innovation.

Table 11.2 Distributed arrangements for exploratory mode

People	Fluid participation in order to incorporate different and changing perspectives
Technologies	Anomalies allow new technologies and new understandings to emerge
Organizational routines	Emphasize exploration and experimentation to promote understandings
Metrics	Emphasize assessments of change and development

The story of Post-It Notes identifies 3M's organizing mode and illustrates how it works. In 1969, while Spencer Silver was conducting experiments in search of a more permanent adhesive, he "stumbled" upon a new substance – glue that did not glue. While Silver admired the structural properties of the new substance under a microscope, he could not imagine how it might be useful:

> "People like myself get excited about looking for new properties in materials. I find that very satisfying, to perturb the structure slightly and just see what happens. I have a hard time talking people into doing that – people who are more highly trained. It's been my experience that people are reluctant just to try, to experiment – just to see what will happen!"
>
> *(Nayak and Ketteringham, 1986: 57–8)*

Silver was obviously not concerned that his new compound might threaten 3M product markets or lead to restrictions on his innovative efforts, but Nayak and Ketteringham (1986: 61) could see such implications: "In this [3M] atmosphere, imagining a piece of paper that eliminates the need for tapes is an almost unthinkable leap into the void."

How does 3M's organizing mode encourage research scientists to uninhibitedly explore technology, focusing exclusively on the knowledge implications? Accounts mention two procedures. The first, the 15 percent rule, requires researchers to spend up to 15 percent of their time on their own projects, exploring ideas unrelated to regular 3M assignments. The 15 percent rule supports not just exploration but "institutionalized rebellion" (Coyne, 1996), which is consistent with another 3M belief: "It is more desirable to ask for forgiveness than for permission" (3M, 1998). 3M's openness to all new findings celebrates the philosophy of William McKnight, 3M's legendary CEO, who said:

> "Those men and women to whom we delegate authority and responsibility . . . are going to want to do their jobs in their own way . . . Mistakes will be made, but if a person is essentially right, the mistakes he or she makes are not as serious in the long run as the mistakes management will make if it is dictatorial and undertakes to tell those under its authority exactly how they must do their job."
>
> *(Coyne, 1996: 10)*

The second procedure involves "bootlegging" – a firm-wide understanding that 3M employees can use any and all of the firm's equipment and facilities for whatever purpose they want during off hours and on weekends. This unrestricted access enables 3M employees to carry out new experiments and develop and explore new ideas. Bootlegging reportedly occurs most often when a product concept lacks official support but committed 3M employees continue to pursue it enthusiastically and passionately.

> "At 3M we've got so many different types of technology operating and so many experts and so much equipment scattered here and there, that we can piece things together when we're starting off. We can go to this place and do 'Step A' on a product, and we can make the adhesive and some of the raw materials here, and do one part over

here, and another part over there, and convert a space there and make a few things that aren't available."

(Nayak and Ketteringham, 1986: 66–7)

The pursuit of bootlegging requires individual "tenacity" because in exploring, researchers often hit dead ends. Yet at 3M a dead end is seen as another positive event:

> "We acknowledge that failure is part of life . . . and we expect failure on a grand scale. For every 1,000 raw ideas, only 100 are written up as formal proposals. Only a fraction of those become new product ventures. And over half of our new product ventures fail. Yet, without these dead ends, there would be no innovation."
>
> *(Coyne, 1996: 12)*

Post-It Notes took 12 years to develop, and often Silver embarked on collaborations designed to draw on knowledge in different domains. New insights can be disorienting, however, as the status of new knowledge is always unknown. Yet in exploration mode, emerging insights are the knowledge drivers that lead to development and change. Silver said: "I've always enjoyed crossing boundaries. I think it's the most exciting part of the discovery process . . . when you bring two very different areas together and find something completely new" (Lindhal, 1988: 16). Another 3M manager described how people experience exploration and how metrics are introduced to guide the process:

> "Among other things, it means living with a paradox: persistent support for a possible idea, but not foolishly overspending because 3M, above all, is a very pragmatic company. It typically works this way: The champion, as his idea moves out of the very conceptual stage and into prototyping, starts to gather a team about him. It grows to say 5 or 6 people. Then, suppose (as is statistically the likely case) the program hits a snag. 3M will likely cut it back quickly, knock some people off the team. But as the mythology suggests, the champion – if he is committed – is encouraged to persist, by himself or perhaps with one co-worker at say a 30 percent or so level of effort. In most cases, 3M has observed that the history of any product is a decade or more long before the market is really ready. So the champion survives the ups and downs. Eventually, often the market does become ripe. His team rebuilds."
>
> *(Peters and Waterman, 1982: 230)*

Exploration takes time. While new insights can provide options for dealing with future issues, it is difficult to identify specific value that has an immediate impact on profits or market share. 3M does not try to identify such impact. Instead, 3M research scientists adopt a long-term perspective and consider how broad product families might benefit from their new ideas. They use developmental milestones to evaluate exploration efforts.

What does 3M tell us about organizing in exploration mode? It suggests that such organizations value experimentation highly because finding explanations for anomalies is a crucial step in knowledge development. People explore to see "what might happen" and encounter "dead ends" that are also valuable in increasing their knowledge.

Table 11.3 Distributed arrangements for normal mode

People	Partitioned roles based on fixed understandings that support normal states
Technologies	With anomalies minimized, technology appears to be stabilized
Organizational routines	Exploit existing understandings to enhance reliable performance
Metrics	Emphasize well-defined and predictable performance

Procedures like the 15 percent rule and bootlegging encourage exploration, and people further enhance their chances of finding new ideas by seeking input from new domains. Evaluation metrics focus on development processes and project milestones, and they do not try to assess the impact on current results.

Organizing in predictable task-performance mode

An alternative but different organizing mode centers attention on efficiency and predictable task performance. In this organizing mode, knowledge is again distributed across the people, technologies, procedures, and metrics of an organization (table 11.3) but the task to be carried out is specific and known and so organizations can assess efficiency and expect predictability.

Hutchins (1995) provides an example with his ethnographic description of how a task, navigating a ship, is carried out. The knowledge needed for this task is distributed across members of the ship's navigation team, the technologies it uses, its routines and procedures, and its assessment processes. While the captain determines the course, midshipmen make sightings, record bearings, time readings, plot the course, and so on. Each person on the team makes a critical contribution. Each action taken is critically determined by the knowledge implicit in the designs of the navigation instruments, the knowledge behind the routines that team members use, and the knowledge reflected in the assessment processes that are the basis of performance evaluations. Team members are required to adhere exactly to the team's routines and procedures because, as bearings are taken, recorded, and analyzed correctly at scheduled, temporally ordered intervals, the team's navigation knowledge is systematically pulled together. Continuously, the captain knows where his ship is and can direct where it should go.

Predictable task performance requires that when people use technologies they must adhere strictly to established procedures and metrics. This approach is consistent with "scientific management" principles and their objective, which is "to increase productivity by streamlining and rationalizing factory operation from cost accounting and supervision to the dullest job on the shop floor" (Zuboff, 1984). Predictable task performance depends fundamentally on the knowledge built into machine designs. As more human skills are transferred to machines and as machines are developed that extend human abilities by incorporating capacities that humans have never had, expanded and predictable task performance becomes realistic and a possibility for continuing future development. This ongoing process has generally reduced the complexity that is dealt with by people, and increased the complexity that is faced by organizations.

> With the growing complexity and size of factories, expanding markets that exerted a strong demand for an increase in the volume of production, and a rising engineering profession, there emerged a new and pressing concern to systematize the administration, control, coordination and planning of factory work.
>
> *(Zuboff, 1984: 41)*

As tasks become more complex, machines incorporate more knowledge to do more things supported by appropriate procedures and metrics. In turn, this opens up new possibilities for what people in organizations can direct and machines can achieve. As the emphasis in this context is on achieving predictable task performance, "faster, better, cheaper" often becomes management's guiding mantra. Given such an organizing mode, experimentation will not be allowed, as it is standardized, repeated operations using known technologies in established ways that enable predictable and productive task performance.

NASA PRIOR TO STS-107

These observations on alternative organizing modes serve as background to help explain some of the challenges that NASA confronted prior to and during the STS-107 disaster. Early in its history, NASA was described as embodying a "can-do" technical culture organized in exploration mode. Vaughan, for example, described NASA as having: "a commitment to research, testing and verification; to in-house technical capability; to hands-on activity; to the acceptance of risk and failure; to open communications; to a belief that NASA was staffed with exceptional people; to attention to detail; and to a 'frontiers of flight' mentality" (1996: 209).

After the success of the *Apollo* program, however, it became increasingly difficult for NASA to obtain resources from Congress (CAIB, 2003). Vaughan (1996) reported that, by the 1980s, the "can-do" culture had given way to a "must-do" culture where the emphasis was on accomplishing more with less. NASA's rhetoric made extravagant claims, suggesting, for example, that NASA was about to achieve "the goal of a 'space bus' that would routinely carry people and equipment back and forth to a yet-to-materialize space station." According to Vaughan, "A business ideology emerged, infusing the culture with the agenda of capitalism, with repeating production cycles, deadlines, and cost and efficiency as primary, as if NASA were a corporate profit seeker" (1996: 210).

This new organizing mode emphasized objective data and predictable task performance. For instance:

> The emphasis was on science-based technology. But science, in FRR [Flight Readiness Review] presentations required numbers. Data analysis that met the strictest standards of

> scientific positivism was required. Observational data, backed by an intuitive argument, were unacceptable in NASA's science-based, positivistic, rule-bound system. Arguments that could not be supported by data did not meet engineering standards and would not pass the adversarial challenges of the FRR process.
>
> ***(Vaughan, 1996: 221)***

Then came the *Challenger* disaster. One of the seven individuals on board, Christa McAuliffe, a schoolteacher, was to "teach schoolchildren from space." A mission that was to celebrate the normality of predictable task performance in space turned tragically into a mission that reminded everyone of the terrors of exploring space. Vaughan (1996) offers several explanations. One draws on Perrow's (1984) theory of "normal accidents" and suggests that, when operating systems have interactively complex and tightly coupled elements, accidents are inevitable. The space shuttle was clearly a complex technology that included 5,396 individual "shuttle hazards," of which 4,222 were categorized as "Criticality 1/1R."[1] The complexity of the interactions that could occur between these identified hazards and some of their interdependent couplings could only increase during such a crisis, making gruesome consequences increasingly likely.

To some extent, however, a predictable task performance organizing mode masks this complexity because, rather than using distributed knowledge to acknowledge and then explore anomalies, procedures often deny or ignore anomalies with the intention of generating at least an appearance of operational predictability (Vaughan, 1996). When developmental technologies are put to use, however, the anomalies they generate eventually have to be dealt with. NASA had responded to anomalous performance by stretching standards and granting waivers, all of which made it possible for the spacecraft to continue to operate, perpetuating an impression of predictable task performance.[2]

Vaughan (1996) noted that, after the *Challenger* accident, the shuttle was no longer considered "operational" in the same sense as a commercial aircraft and that NASA continued to combine organizing modes. In 1992, for example, Daniel Goldin, the incoming NASA Administrator, insisted that a reorganized NASA would do things faster, better, and cheaper without sacrificing safety. Goldin's approach implies that at the highest levels of NASA there was a lot of emphasis placed on increasing efficiency and ensuring predictable task performance (CAIB, 2003: 103). Similarly in 2001, Goldin's successor, Sean O'Keefe, tied future Congressional funding to NASA's delivery of reliable and predictable shuttle flight performance in support of the International Space Station. Such a promise implied a high ability to maintain predictable task performance – similar to what one might achieve, for example, in managing assembly-line production.

FLIGHT STS-107

Strapped for resources, committed to maintaining a demanding launch schedule in support of the International Space Station that many managers knew could not be met, and with parts of itself committed to different organizing modes, NASA faced the series of events that emerged over the 14-day flight period of flight STS-107 and culminated in the destruction of the shuttle and its crew. Figure 11.1 charts how, for the first nine days of the 14-day flight, events unfolded and interactions occurred between different NASA groups, and the various technologies, procedures, metrics, and tools that comprised the elements of NASA's distributed knowledge system. Figure 11.1 is not intended as a complete set of events and responses that occurred over this period. Rather, it is a set chosen to illustrate and summarize the sequence of events mentioned in our narrative that, in turn, highlights how indeterminacy came to dominate NASA's distributed knowledge system.

Foam shedding

The most critical event for STS-107 occurred around 82 seconds into the flight. One large object and two smaller objects fell from the left bipod area of the external tank and the large object then struck the underside of the orbiter's left wing. Although videos and high-speed cameras captured this event, the image resolution was fuzzy. Only when images with better resolution became available the next day did it actually become evident that something significant could have happened. The new images were also not completely clear and so questions continued to surround the event.

The intention of the original shuttle design was to preclude foam shedding and make sure that debris could not fall off and damage the orbiter and its thermal protection system (TPS). On many shuttle flights, however, objects have fallen from the external tank and some have hit the orbiter's left wing damaging parts of the TPS. Over time NASA has learned that, by making repairs after each flight, it can take care of this damage.

Categorizing foam shedding

On several shuttle flights, chunks of foam have fallen from the external tank's forward bipod attachment as occurred on STS-107. Early on, these events were classified as "in-flight anomalies" that had to be resolved before the next shuttle could launch. After the orbiter was repaired following a foam-loss event in 1992, however, it was determined that foam shedding during ascent did not constitute a flight or safety issue, and so the assessment metrics changed. For many managers, foam loss became an "acceptable risk" rather than a reason to stop a launch. Nevertheless, the foam shedding that occurred during the launch of flight STS-107 attracted the attention of the Intercenter Photo Working Group. Concerned with the size of the debris and the speed at which it was moving, the Intercenter Photo Working Group anticipated a wider investigation that would explore the incident and so it requested additional on-orbit photographs from the Department of Defense.

The Intercenter Photo Working Group classified the STS-107 foam loss as an "out-of-family event,"[3] meaning it was something unusual that they had not experienced

	Prior to incident	1/16	1/17	1/18, 19	1/20	1/21	1/22	1/23	1/24	2/01	
ARTIFACTS											
	5396 documented shuttle hazards		Strike is classified as out of family				Critical day for rescue of crew passes unnoticed				Indeterminate event due to interactive complexity
Foam shedding	Occurs, though design prohibits: a turnaround issue	Huge block of foam breaks off and hits left wing									
Tiles	Consistently damaged in flights: a turnaround issue	Extensive damage								Collapses on re-entry	
RCC panels	Have little ability to withstand kinetic energy. Foam loss is believed not to threaten panels	Panels extensively damaged								Thermal Protection System (TPS) collapses on re-entry	
Shuttle	No crew escape system	Compromised								Destroyed	
PEOPLE/GROUPS											
	Foam shedding, etc. is inevitable, "normal," acceptable risk										
Mission Control	Based on 112 past flights, foam loss is categorized as being a not-safety-of-flight issue	Ham asks for implication for STS-114, the next scheduled shuttle flight	Registered belief that strike posed no safety issue. No rush. Appoint Debris Assessment Team	Did not believe strike posed a safety issue. No rush – analysis Did not answer Mission Action request	Initial attempt is made to classify foam strike as within experience base	Meet DAT Ham explores rationale used for flying, and implications of STS-107 issues for later flights	Ham asks if extra time to get left wing imagery will cause delays Cancels requests for DOD imagery. Know strike is a safety issue but do not follow up	System is tipping towards view that there is no problem and an emphasis on predictable performance	Concludes after DAT team presentation that foam strike is a "turnaround" issue	Arranged post-flight photos of STS-107 to consider implications for STS-114	The foam strike was seen as more of a threat to the shuttle schedule rather than a direct threat to STS-107
Intercenter photo working group	Wants foam losses on earlier flights classified as in-flight anomalies	Reviewed imagery. Did not initially see strike because images unclear	Now see strike Seek more imagery to clarify size. Email imagery								Action requires approval of Mission Management
Debris Assessment Team (DAT)	Procedures to be developed after a team is appointed to enable a report by a given date		Appointed to analyze foam strike	An engineer from Boeing works on analysis over weekend using Crater model	Expanded group size and expertise. Want more images of left wing	Meeting 1 to assess damage Initiate request for images of shuttle in orbit	Meeting 2: Need to but can't prove the need for imagery in support of imagery request A Catch-22	Meeting 3: Still want imagery as they believe the foam block was the size of a large cooler going 500 mph. Cannot get the need for imagery through	Presentation to Mission Control and engineers emphasized their uncertainty about where the debris struck. Could not show for certain it was a safety-of-flight issue		Because it is not designated a "Tiger Team," DAT issues are not given priority by mission control managers
Program Requirements Control Board	Declared bipod foam loss an "action," not an "in-flight anomaly" on previous flight										Lowered urgency of action in response to foam-shedding event
Rodney Rocha, chief engineer, Thermal Protection System	Coordinates NASA engineering resources and works with contract engineers at United Space Alliance (USA)		Coordinator of DAT with Pam Madera of USA	Made Mission Action Request that *Columbia* crew inspect left wing		Requests JSC Engineering for imagery of vehicle on-orbit	Believes request for imagery has been denied. Unsent email expresses deep concerns	Disagrees with Schomberg that foam impact is in experience base			Routing request via JSC Eng. reduces its salience. Rocha's concerns remain unresolved
METRICS/TOOLS											
	Foam shedding assessed after flights. No metrics to assess it during flights										Never knew if there was a flight safety issue or not
Cameras	On earth, satellites, planes as well as shuttle but abilities have atrophied due to budget cuts	Picture unclear or missed event	Backups haven't worked	Need to get supplementary images							Cameras unreliable and provide unclear pictures for analysis
Crater model	Assumes small-sized debris, model gives quick fix on TPS penetration depth			First use of model for ongoing mission	Use of model confirms need for imagery						Results from use of tool are inconclusive
Int'l Space Station schedule	Importance of core complete date Feb. 19, 2004 linked to Sean O'Keefe										Scheduled launches, already behind, will be backed further
ESTABLIHED ROUTINES/PROCEDURES											
	Debug through categorization waivers allows adherence to schedule. Delays only permitted by convincing and objective data		Analysis team needs "Tiger" status to get Mission management's attention			Request for imagery not done through correct formal channel		Email directs that future DOD request should go through correct formal channels			Routing of image requests outside regular chain of command leads to unclear importance
Flight Readiness Review	Foam shedding is made an accepted risk not a safety of flight issue										
Tiger-Team Process	Detailed procedures and checklists to be explored given out-of-family event		Not designated								Never activitated
	Prior to incident	Day 1	Days 2 & 3	Day 4	Day 5	Day 6	Day 7	Day 8	Day 9	Day 17	

Figure 11.1 Partial response graph of STS-107 disaster

before and they wanted to obtain more data to explore it further. As it was not clear to the group whether or not it posed a "safety-of-flight issue,"[4] they asked for foam-loss events on earlier flights to be classified as in-flight anomalies so that they could get historical data to examine and explore previous events. Tracing out the classifications used to sort, order, and distinguish everyday events helps an organization to understand the events it faces and their significance (Bowker and Star, 1999). The outcomes of such sorting processes manifest themselves not only in event classifications but also in metrics used to evaluate events. They are a part of the organizational knowledge that accumulates over time. They not only determine appropriate responses to events but they also contribute to the evaluation criteria used to set the stage for new rounds of classification that will guide responses to emerging events faced on future flights.

However, other groups in NASA classified the foam-loss event observed during the launch of STS-107 in different ways. These different classifications were not simply semantic exercises. Rather, different metrics led to different ways of classifying the anomaly and, in turn, such an evaluative classification established a critical link to the procedures that NASA would use in response to this and other emerging events. The CAIB report (2003: 121–74) suggests that, from the beginning, those in Mission Control interpreted the foam-loss event as being in the "in-family" category, a "turn-around" rather than a "safety-in-flight" issue. They were often most concerned about how actions and delays in dealing with issues arising from flight STS-107 might reduce NASA's ability to meet downstream schedules and commitments.

The CAIB report also states that other groups in NASA classified the foam-loss event as an "out-of-family" event. These included the Intercenter Photo Working Group and several of Boeing's engineering analysts. The latter described the event as "something the size of a large cooler [that] had hit the Orbiter at 500 miles per hour." These groups wanted more information about the foam strike, specifically on-orbit photographs, so that they could check and determine by direct observation whether or not the event posed a safety of flight issue.

Deciding on a response

An organization with institutionalized processes for the identification of out-of-family events is reflexive enough to realize that the significance of emergent events is not always known. Given the recognition that organization members might identify an event as "out-of-family," NASA had prepared pre-programmed responses for handling such an occurrence. If an emergent event was classified as "out-of-family," its procedures required the automatic formation of an assessment team to lead an in-depth examination of the event.

Depending on exactly how the event was classified, however, an assessment team could be appointed at different status levels. For incidents that were out-of-family, the routine stipulated the appointment of a "Tiger Team" that would have wide and extensive authority to ask questions and get things done in order to find out quickly what had occurred. For incidents that were in-family, in contrast, an assessment team with less authority could be appointed with the expectation that it could take more time to work out what had occurred and that their report could be made on a specific but less urgent date.

In the case of STS-107, some groups involved in NASA's distributed knowledge system classified the falling foam as an out-of-family event while others, including Mission Control, thought it should be classified as an in-family event. Their preferences reflected the different metrics that were part of their different organizing modes. As Mission Control had ultimate overall charge, they made a unilateral decision to appoint a lower-status Debris Assessment Team (DAT), as was consistent with their assessment that they were dealing with an in-family event. This was even though another part of NASA – the Intercenter Photo Working Group – had already classified the event as out-of-family and, according to NASA's procedures, a reclassification that leads to a different status for the investigating team should not occur. As the DAT did not have "Tiger Team" status, it had no authority to carry out the preprogrammed actions and checklists that become a part of Mission Control's procedures when a Tiger Team is appointed. "This left the Debris Assessment Team in a kind of organizational limbo, with no guidance except the date by which Program managers expected to hear their results: January 24th" (CAIB, 2003: 142).

Developing a response

As there were no in-place procedures to help DAT build up its knowledge of what had occurred in the foam strike, DAT improvised, seeking insights wherever it could find them. For example, the group used a mathematical modeling tool called "Crater" to assess what the damage to the wing might be. While the use of this tool was inappropriate in that it was calibrated to assess damage caused by debris of a much smaller size than that which was estimated to have hit the underside of *Columbia*'s wing (1 cubic inch vs. over 600 cubic inches), it was one of the few tools available. It was also the first time that the tool had been used to assess damage to a flight that was still in orbit. Crater model analyses were designed to be conservative, i.e., given alternative statistical errors the analyses minimized the likelihood of failing to identify a "safety in flight" problem.

The prediction from the Crater model analysis was that the debris might well have compromised the underside of the wing, generating an "out-of-family" event that on re-entry could expose the shuttle to extremely high temperatures. However, having obtained a disturbing assessment, the DAT team then sought to discount the prediction. They reasoned that, as the tool was conservative, its design was to avoid missing a safety in flight issue. Statistically, then, the model might be making a false-positive identification of a safety in flight problem. Second, they reasoned that the Crater model did not factor in the additional padding that STS-107 had packed on the underside of its wing. To explore the effectiveness of this additional padding, however, the engineers had to determine the location and angle of the foam strike impact. A "transport" analysis surfaced a scenario whereby the foam strike likely hit at an angle of 21 degrees. Would such a strike compromise the reinforced carbon

carbon (RCC) coating the wing? To answer this question, DAT resorted to another mathematical model calibrated to assess the impact of falling ice. It predicted that strike angles greater than 15 degrees would result in RCC penetration and portend disaster. However, as foam is not as dense as ice, DAT again decided to adjust the analysis standards. These steps led it to conclude that a foam strike impact angle of up to the suspected 21 degrees might not have penetrated the RCC. The CAIB report states: "Although some engineers were uncomfortable with this extrapolation, no other analyses were performed to assess RCC damage" (2003: 145).

In addition to the on-orbit photographs that had been requested by the Intercenter Photo Working Group, DAT initiated its own requests for on-orbit photographs. Rather than following the chain of command through Mission Control, however, DAT requested the imagery through an Engineering Directorate at Johnson Space Center, a group that DAT's leader was familiar with. The fact that the request for on-orbit photographs was made by an engineering unit rather than by Mission Control's flight dynamics officer signaled to Ham, the STS-107 flight controller, that the request was related to a non-critical engineering need rather than to a critical operational concern. Acting for Mission Control, therefore, she canceled the request to the Department of Defense for on-orbit images. Ham's action terminated the requests for on-orbit photographs made by both the Intercenter Photo Working Group and DAT (CAIB, 2003: 153).

DAT did not realize that, in canceling the request, Ham had not known that it was a DAT request. She reported that she terminated the request based simply on its source that was not DAT. Although DAT members did not know about this confusion, they did know that Ham had turned down their request. They also knew that, from the start, Mission Control had wanted to classify the foam strike event as "in-family," a classification that would also nullify their request for on-orbit photos. DAT, therefore, was put in the "unenviable position of wanting images to more accurately assess damage while simultaneously needing to prove to Program managers, as a result of their assessment, that there was a need for images in the first place" (CAIB, 2003: 157). In other words, the DAT team felt caught in a Catch-22 bind where they were required to show objective evidence for a safety in flight issue, i.e., on-orbit photos, but in order to get this evidence they first had to provide other convincing evidence demonstrating that there was a safety in flight issue.

Engineers and DAT members continued to be concerned. A structural engineer in the Mechanical, Maintenance, Arm and Crew Systems, for example, sent an email to a flight dynamics engineer stating: "There is lots of speculation as to extent of the damage, and we could get a burn through into the wheel well upon entry." The engineer leading DAT, Rodney Rocha, wrote an email that he did not send because "he did not want to jump the chain of command":

> In my humble technical opinion, this is the wrong (and bordering on irresponsible) answer from the SSP and Orbiter not to request additional imaging help from any outside source . . . The engineering team will admit it might not achieve definitive high confidence answers without additional images, but, without action to request help to clarify the damage visually, we will guarantee it will not.
>
> *(CAIB, 2003: 157)*

Despite the fact that parts of the distributed knowledge system wanted to find out more about the damage the foam strike had caused, the overall system was prematurely stopping exploration efforts and tipping toward an organizing mode focused on achieving predictable task performance. For instance, after Ham canceled the request for imagery to the Department of Defense and still a day before DAT's presentation of its findings to Mission Control, a NASA liaison to USSTRATCOM had already concluded what the results would be and sent this email:

> Let me assure you that, as of yesterday afternoon, the Shuttle was in excellent shape, mission objectives were being performed, and that there were no major debris system problems identified. The request that you received was based on a piece of debris, most likely ice or insulation from the ET, that came off shortly after launch and hit the underside of the vehicle. Even though this is not a common occurrence it is something that has happened before and is not considered to be a major problem.
>
> *(CAIB, 2003: 159)*

In an effort to ensure predictable task performance, the email continued:

> The one problem that this has identified is the need for some additional coordination within NASA to assure that when a request is made it is done through the official channels . . Procedures have been long established that identify the Flight Dynamics Officer (for the Shuttle) and the Trajectory Operations Officer (for the International Space Station) as the POCs to work these issues with the personnel in Cheyenne Mountain. One of the primary purposes for this chain is to make sure that requests like this one do not slip through the system and spin the community up about potential problems that have not been fully vetted through the proper channels.
>
> *(CAIB, 2003: 159)*

The response

On Day 9, DAT made its formal presentation to Mission Control and a standing-room-only audience of NASA engineers lined up to hear stretched out into the hallway. The DAT members had worked in exploration mode and had wanted to provide evidence clearly identifying what had occurred. In their presentation, DAT members stressed the many uncertainties that had plagued their analyses – for instance there were still many uncertainties as to where the debris had hit and there were many questions stemming from their use of the Crater model. Because of these uncertainties, they could not prove that there was a definite safety of flight issue. The Mission Control Team then saw no reason to change its view, held all along, that the foam strike was not a safety of flight issue. According to the CAIB report, however, "engineers who attended this briefing indicated a belief that management focused on the answer – that analysis proved there was no safety-of-flight issue – rather than concerns about the large uncertainties that may have undermined the analysis that provided that answer" (2003: 160).

Epilogue

After a series of interactions between the different parts of NASA's distributed knowledge system, what began as an "out-of-family" event was eventually categorized as an "accepted risk" and no longer a "safety of flight" issue. As the shuttle re-entered the Earth's atmosphere on February 1, hot plasma breached the RCC tiles and the shuttle disintegrated. The CAIB report said:

> Management decisions made during *Columbia*'s final flight reflect missed opportunities, blocked or ineffective communications channels, flawed analysis, and ineffective leadership. Perhaps most striking is the fact that management – including Shuttle Program, Mission Management Team, Mission Evaluation Room, and Flight Director and Mission Control – displayed no interest in understanding a problem and its implications. Because managers failed to avail themselves of the wide range of expertise and opinion necessary to achieve the best answer to the debris strike question – "Was this a safety-of-flight concern?" – some Space Shuttle Program managers failed to fulfill the implicit contract to do whatever is possible to ensure the safety of the crew. In fact, their management techniques unknowingly imposed barriers that kept at bay both engineering concerns and dissenting views, and ultimately helped create "blind spots" that prevented them from seeing the danger the foam strike posed.
>
> *(CAIB, 2003: 170)*

While evidence supports this conclusion, it was also possible in real time that a path to explore the out-of-family event could have been pursued rather than closed off, for core parts of NASA are not just technically driven but they are also organized in exploration mode. There was also a period during the interactions between NASA's different groups when what conclusion would be reached about the foam strike was also indeterminate, i.e., the distributed knowledge system could have tipped either way. How did the system tip toward a predictable task-performance mode and prematurely close off exploration?

One speculative proposition might be that time pressure played a role. There were in fact many sources of time pressure influencing this situation and one related to the time-dependent options for a possible rescue. Specifically, if it had been determined by Flight Day 7 that STS-107 had suffered catastrophic damage, a rescue mission could have been launched with the *Atlantis* shuttle. It is not clear, however, that Mission Control members seriously thought about this deadline or even considered a rescue. Specific quotes suggest that, although the damage that had been caused by the foam shedding was not known, at least some Mission Control members simply did not consider a rescue to be either necessary or an option, and others did not consider it even to be possible.

Another source of time pressure relates to the commitments that had been made by NASA leaders to service the International Space Station. Although many NASA leaders realized that the completed phase (Node 2) for the International Space Station planned for February 19, 2004 would not be achieved, they were nevertheless committed to accomplishing as much as they could toward this goal. This commitment almost certainly encouraged organizing in support of predictable task-performance mode rather than exploration mode. The efforts of top leaders to fulfill performance commitments may also so pervade a resource-strapped organization that it becomes difficult for any group within to counter the implicit lack of support for the exploration efforts that it would like to undertake.

To appreciate NASA's difficulties in dealing with these pressures, however, one must also consider the context that management faced. Foam shedding and potential tile damage was just one of over 5,396 known and documented hazards associated with the shuttle and it was not the problem accorded the highest concern. It's not difficult to imagine that NASA's top management believed that a significant deployment of resources to address and prevent damage stemming from these thousands of identified hazards would permanently stall its programs, making it impossible to fulfill commitments to Congress that had been made with respect to the International Space Station or anything else.

DISCUSSION AND CONCLUSION

Echoing our idea that NASA incorporates dual organizing modes, Ron Dittemore, the space shuttle project manager, testified on March 6, 2003:

> "I think we're in a mixture of R&D and operations. We like to say that we're operating the fleet of Shuttles. In a sense we are, because we have a process that turns the crank and we're able to design missions, load payloads into a cargo bay, conduct missions in an operating sense with crew members who are trained, flight controllers who monitor people in the ground processing arena who process. In that sense we can call that operations because it is repeatable and it's fairly structured and its function is well known. The R&D side of this is that we're flying vehicles – we're blazing a new trail because we're flying vehicles that are, I would say, getting more experienced. They're getting a number of flights on them, and they're being reused. Hardware is being subjected over and over again to the similar environments. So you have to be very careful to understand whether or not there are effects from reusing these vehicles – back to materials, back to structure, back to subsystems."
>
> *(CAIB, 2003: 20)*

Although organizations are responsive systems, events can only be understood and acted upon from certain perspectives that take account of particular situations. Early on, there are possibilities for one of many organizational response patterns to be activated. As sequences of events unfold, they begin generating overall constraints on the paths that an organization may pursue in its responses. Such constraints spread through the system over time. Beyond a critical threshold, such processes start to tip a system of distributed elements into one or other organizing modality where knowledge and actions start to become consistent with a particular perspective.

Timing and temporality are vital in this process. The sequences in which specific elements of a response system are activated become critically important in determining the overall pattern of responses that becomes accepted in an organization. The impact of doubts and ambiguities, for instance, depends upon not only who raises them but also when they are raised. Similarly, the order in which alternative metrics, artifacts, and routines are activated critically impacts patterns of organizational response.

It is often difficult if not impossible for the people in an organization to determine the status of an event in real time. This is because many organizations are often operating in multiple modes. Events can be understood and acted upon, however, only given certain perspectives. To the extent that an organization is operating in a dual mode, critical ambiguities emerge that prevent it from generating shared "collateral experiences" (March et al., 1991). Any of the four elements that make up a complex response system can help generate indeterminacy and, as a result, an organization may be in no position to initiate any response even if it has been designed for high reliability (Weick and Roberts, 1993).

In the case of STS-107, all elements functioned to generate indeterminacy that had to be, but could not be, resolved within a short time frame. Despite access to data, vast resources, and widespread goodwill, NASA could not in real time identify the significance of the foam-loss event and the consequent emerging crisis, and it took a heavy toll. Indeterminacy of data is one explanation for NASA's inability to act in real time.

These observations and processes have implications for learning in other organizations. In predictable task-performance mode, management stabilizes in order to ensure predictable performance. Learning according to this perspective implies an ability to do existing tasks progressively better through a process of learning by doing. In exploratory mode, in contrast, management allows ideas and architectures to emerge and supports discovery and creativity. Learning according to this perspective implies discovering unknown processes for building new things and accomplishing new aims.

In seeking to incorporate both, NASA may have placed itself somehow in between both. On the one hand, the pressures that NASA is under and to some extent has selected require it to stabilize and rationalize its activity. But as NASA also necessarily deals with technologies that have emergent and changing designs, learning that is appropriate for predictable task performance may actually block the learning that is needed to support developing technologies.[5] As NASA is operating in two modes, simultaneously, processes generating indeterminacies may be inherent in the "One NASA" vision.

ACKNOWLEDGMENTS

We thank participants at the NSF conference on design held at NYU for their valuable comments. We thank Moshe Farjoun and Bill Starbuck for their comments and help.

NOTES

1 CRIT 1/1R component failures are defined as those that will result in loss of the orbiter and crew.

2 Of the 4,222 that are termed "Criticality 1/1R," 3,233 have waivers. Waivers are granted whenever a Critical Item List component cannot be redesigned or replaced. More than 36 percent of these waivers have not been reviewed in 10 years.

3 An out-of-family event is defined as: "Operation or performance outside the expected performance range for a given parameter or which has not previously been experienced." An in-family event is "A reportable problem that was previously experienced, analyzed, and understood. Out of limits performance or discrepancies that have been previously experienced may be considered as in-family when specifically approved by Space Shuttle Program or design project."

4 No safety of flight issue: the threat associated with specific circumstance is known and understood and does pose a threat to the crew and/or vehicle.

5 In their analysis of scientific and technological challenges associated with the construction of aircraft capable of attaining satellite speeds, Gibbons et al. (1994: 20–1) point out that "discovery in the context of application" generates fundamental discontinuities with previous experiences.

REFERENCES

Bowker, G.C., and Star, S.L. 1999. *Sorting Things Out: Classification and its Consequences*. MIT Press, Cambridge, MA.

CAIB (*Columbia* Accident Investigation Board). 2003. Report. Government Printing Office, Washington, DC. 6 vols.: vol. 1. www.caib.us/news/report/volume1/; vol. 6: www.caib.us/news/report/volume6/.

Callon, M. 1998. The embeddedness of economic markets in economics. In M. Callon (ed.), *The Laws of the Markets*. Oxford: Blackwell, pp. 1–57.

Coyne, W.E. 1996. *Building a Tradition of Innovation*, 5th UK Innovation Lecture. Department of Trade and Industry, London.

Garud, R., and Rappa, M. 1994. A socio-cognitive model of technology evolution. *Organization Science* 5(3), 344–62.

Gibbons, M., Limoges, C., Nowotny, H., Schwartzman, S., and Trow, M. 1994. *The Production of Knowledge: The Dynamics of Science and Research in Contemporary Societies*. Sage Publications, New Delhi.

Hutchins, E. 1995. *Cognition in the Wild*. MIT Press, Cambridge, MA.

Lindhal, L. 1988. Spence Silver: a scholar and a gentleman. *3M Today* 15, 12–17.

March, J.G., Sproull, L.S. and Tamuz, M. 1991. Learning from samples of one or fewer. *Organization Science* 2(1), 1–13.

Nayak, P.R., and Ketteringham, J.M. 1986. *Breakthroughs!* Rawson Associates, New York.

Perrow, C. 1984. *Normal Accidents: Living with High-Risk Technologies*. Basic Books, New York.

Peters T.J., and Waterman, R.H. 1982. *In Search of Excellence*. Harper & Row, New York.

3M. 1998. *Innovation Chronicles*. 3M General Offices, St. Paul, MN.

Tsoukas, H. 1996. The firm as a distributed knowledge system: a constructionist approach. *Strategic Management Journal* 17(Winter), 11–25.

Weick, K.E., and Roberts, K.H. 1993. Collective mind in organizations: heedful interrelating on flight decks. *Administrative Science Quarterly*, 38(3), 357–81.

Vaughan, D. 1996. *The Challenger Launch Decision: Risky Technology, Culture, and Deviance at NASA*. University of Illinois Press, Chicago, IL.

Zuboff, S. 1984. *In The Age of the Smart Machine*. Basic Books, New York.