Submission Number 14982

DATA INDETERMINACY IN ORGANIZATIONS

Roger Dunbar[**]
Stern School of Business
New York University
New York NY 10012
rdunbar@stern.nyu.edu
&
Raghu Garud
Smeal College of Business
The Pennsylvania State University
University Park, PA, 16802
rgarud@psu.edu

March 2007

[**] 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.

DATA INDETERMINACY IN ORGANIZATIONS

Abstract

The functioning of organizations presupposes an ability to make sense of emergent situations. Not only should such ability enable an organization to respond to emerging situations in real-time, but it should also enable an organization to learn over time.

By focusing on how organizations respond to emerging situations, however, we encounter an interesting puzzle. What if a situation is recognized as emerging but its nature remains indeterminate? A situation is indeterminate if multiple perspectives within an organization attribute different meanings to it, thereby generating ambiguities that confuse its overall significance.

The idea that situations may be indeterminate throws a new light on knowing and learning in organizations. If emerging situations are indeterminate in real-time, how can organizations respond to them? Moreover, how can organizations accumulate knowledge over time as they are faced with indeterminate situations?

On the early morning of February 1, 2003, the incoming Columbia Shuttle Flight 107 started reentry 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 reenter the earth's atmosphere, was traveling at around Mach 23, and its wing temperatures were set to rise to over 2800 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 cu. in. 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 reentered the earth's atmosphere 16 days later, hot plasma gas flowed into the spacecraft setting off chain reactions that destroyed the shuttle's systems and crew.

A report compiled by a board of eminent evaluators (The Columbia Accident Investigation Report – CAIB) came to the following conclusion:

> "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." (The CAIB Report, 2003:170)

Such an attribution of faulty judgment to mission control managers is consistent with a "decision rationality" perspective (Brunsson, 1985) that suggests that managers should evaluate emergent situations, weigh the pros and cons and initiate appropriate actions. From this perspective, decisions go awry due to the cognitive biases (or blind spots as the CAIB Report suggests) of decision-makers (c.f. Kahneman and Tversky, 1972). However, an alternative perspective suggests that organizational responses emerge not because of the decision-making abilities of individual managers, but because various elements that lie distributed across an organization are activated to identify the situation and select an appropriate response (c.f. Hutchins, 1995; Callon, 1998). These activated elements include not only the skills and competencies of different people, but also the routines an organization uses to identify situations and direct its responses, the technologies the organization relies on to carry out such routines, and the evaluation metrics that the organization and its different units use to assess situations.

From such an “action rationality” perspective (Brunsson, 1985), organizations make sense of emergent situations only by activating and then coordinating these distributed elements. A specific organizational response also depends upon the ways in which the elements address different operating modes. For instance, the various distributed elements could be designed to support ongoing task-performance, or they could be designed to support the exploration of new possibilities. In either case, distributed organizational elements facilitate the sorting and ordering of information about emergent situations into categories to generate an appropriate organizational response consistent with a designed, operating mode (Bowker and Star, 1999).

Difficulties in sorting and ordering processes arise when organizations are designed to deal with multiple operating modes. Emergent situations can then potentially be categorized and acted upon in multiple ways. Under these circumstances, organizations cannot simply quickly classify events and situations as belonging to one category or another. Instead, a struggle may ensue within the organization to classify a situation in order to generate an organizational response. If there is not enough time for these identification and classification battles to play out, the situation may remain indeterminate and, therefore, outside of the organization’s influence. This is the basic proposition that we address in this paper.

ORGANIZATIONS AND DISTRIBUTED RESPONSE

Organization theorists have long considered organizations as responsive systems. One position is that responses are driven by a logic of consequences, i.e., “actions are chosen by evaluating their probable consequences for the preferences of the actor” (March and Simon, 1993:8). From this perspective, managers can and ought to evaluate emergent situations, weigh pros and cons and, eventually, decide upon an appropriate course of action.

March and Simon (1993) proposed a second logic – one of appropriateness. As they stated, “Actions are chosen by recognizing a situation as being of a familiar, frequently encountered type and matching the recognized situation to a set of rules” (March and Simon, 1993:8) They noted in the second edition of their book that they would place relatively less emphasis on “analytically rational as opposed to rule-based action” (March and Simon, 1993: 5).

The power of organizations, then, lies not just in inducing cooperation between calculating agents, but also in embedding them within a set of rules, technologies and evaluation metrics that shape understandings and responses to emergent situations (Garud and Rappa, 1994; Hutchins, 1995; Callon, 1998; Beunza and Stark, 2004). Calculation may unfold, but within a network of distributed elements appropriately organized according to a particular operating mode. Oftentimes, the calculative process is so deeply embedded in such networks and the associated operating mode that they appear to be almost intuitive.

-- Table 1a here --

To generate a response to emergent situations, the elements that contain an organization’s distributed knowledge must be activated (Tsoukas, 1996). For instance, the evaluation of a potential strategic opportunity occurs as people with specific skills become involved and as routines, technologies and metrics are activated. In other words, it is only through the activation of these distributed elements that an organization makes sense of emergent situations and simultaneously constructs a response – a response process that reflects what Brunsson (1985) labels as “action rationality.” Consequently, an organization’s response pattern will depend upon the ways in which the multiple elements have been designed and then on how they operate together.

Hutchins (1995) ethnographic account of how a ship navigates a way into a harbor illustrates distributed cognition in action. The captain and his team are clearly important. But equally important are the midshipmen carrying out navigating routines including those making sightings, those recording bearings and those timing the readings, etc. In turn, the actions and judgments of these people are profoundly shaped by the knowledge implicit in the designs of the instruments they use and in the evaluation metrics that they employ to assess the readings their routines generate. The rules for taking a bearing at regularly scheduled intervals ensure that these distributed elements taking action in a temporally ordered manner continually ascertain the position of the ship.

Hutchins (1995) study of a ship navigating into port focuses on a single objective and an operating mode consistent with this objective. Each step in the process makes intuitive and obvious sense. It is not unusual, however, for organizations to embrace multiple operating modes simultaneously (Brunsson, 1985). They do so because they have to relate to and cater to multiple social groups, each with its own set of objectives, as was the case, for example, with NASA. To generate legitimacy and resources, NASA began catering to different constituencies, and, as a result, incorporated into its organizational processes at least two different sets of objectives and associated operating modes. One was a quest for safe, regular shuttle flight operations. A second was a quest for exploration – “to go where no one has been before.”

These two goals place contradictory demands on organizations. For instance, in organizations designed primarily to enhance task-performance, people may be socialized to enhance understandings that support normal states, technologies are designed to black box their functioning, organizational routines are designed to exploit existing understandings to enhance reliable performance, and metrics emphasize well-defined and predictable task performance (Table 1b). In contrast, in organizations designed primarily to enhance exploration, people may be socialized to enhance fluid participation, technologies are designed to be malleable and transparent, routines are designed to emphasize experimentation, and metrics may emphasize change and development and the value of the unexpected (Table 1c).

-- Table 1b and 1c here --

Under these conditions, as organizations encounter situations and generate data from those situations, they are likely to find that, in real time, the meaning of the data generated to the organization is indeterminate. By this we mean that the status of the events that unfold – whether they should be treated as a part of a system that promotes experimentation, or as a part of a system that promotes task-performance – is not clear. Such organizations struggle to classify the data from emergent situations in real time in order to generate an appropriate response. But classification using inconsistent criteria is ultimately impossible and, in many instances, the time available to the organization to classify a situation runs out.

To establish this proposition, we deconstruct the processes and actions that impacted the ill-fated STS-107 shuttle (the Columbia) flight. 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 appropriate for achieving predictable task performance can be directly at odds with understandings appropriate for exploration. We show that as individuals operating in real-time attempt to accommodate both organizing perspectives simultaneously within a distributed response system, the significance of available real-time data becomes indeterminate to those charged with making decisions as ways to react or respond become impossible for these individuals to discern. In situations that demand high-reliability because of threats to human life, the emergence of indeterminacy can have disastrous consequences, as was the case with STS-107.

TECHNO-ETHNOGRPAHY

How can one study and write about an unfolding situation from a distributed perspective? The method should facilitate an ability to record in real time the different perspectives that emerge and the activation of the different organizational elements in this emerging process. Hutchins (1995) addressed these challenges by deploying recording devices at various places of a ship as it steered into harbor. He then analyzed the recordings made by these technologies to piece together the activities that unfolded to establish his perspective on distributed cognition.

We are the beneficiaries of such “techno ethnography” in our study. In compiling a detailed 300-page report, the Columbia Accident Investigation Board members (2003) certainly relied on traditional approaches such as interviews and notes to generate a narrative of the situation as it unfolded. But equally importantly, the CAIB board members relied on the data generated and automatically recorded in the technologies, routines and metrics activated by NASA during space flights. Like a “black box” of a plane, these distributed organization elements aided by new information technologies generated observations such as photographs, emails and the like at different stages of the flight of STS-107. As a result, these data offered board members an opportunity to reconstruct the emerging situation as if it was occurring in real time.

NASA IN TWO MODES

To appreciate how normal and experimental modes were institutionalized into NASA, we briefly examine the agency’s history. In the 1960s, President Kennedy established a clear vision when he said that, through NASA, the US would establish technical superiority over the Soviet Union in space and this achievement would be symbolized with the landing of the first man on the moon within a decade. With generous government funding, a clear mission, and a “can-do” attitude that required the development of many new technologies, and decentralized facilities carrying out the many different projects that a flight to the moon required, NASA organized to achieve its mission (Vaughn, 1996).

NASA’s engineers worked in different facilities on separate projects where their experiments uncovered new problems and the need for further experimentation, and budget overruns became common. It was soon evident that NASA’s overall mission progress was extremely slow. Unless progress somehow speeded up, NASA would not be able to put a man on the moon by the end of the decade and as this became public knowledge, it would lose its base of political support. In 1963, George Mueller joined NASA as director of the Office of Manned Space Flight (OMSF) and he immediately set about integrating ongoing efforts and speeding up development progress. To integrate efforts, he introduced “system engineering”, a set of project management techniques to codify good scientific, engineering and management practices to control the design projects at the various NASA centers and centrally integrate them into the overall mission that top management could track (Johnson, 2002).

To speed up progress, Mueller imposed “all-up testing” processes on NASA. To this point, NASA’s developmental process had emphasized learning by failing: engineers would test a part, it would initially fail, and they would keep on testing until it eventually worked satisfactorily. They would then combine parts into components and test the new combination and after it too failed, they would make modifications and eventually, it too would work perfectly. This was a time-consuming process and "all-up testing" was to reduce the learning-by-failing cycle. Under Mueller's new regime, engineers had to design and fabricate components correctly so there would be no failures. If the design and fabrication were done correctly, the component would work and so rather than testing each component separately, they would all be tested together (Murray and Cox, 1989).

This speeded up development and saved costs. It also introduced unknown risks due to the unknown interactions that would occur between the untested, newly combined components. Implicitly, NASA was accepting the risk that any space flight could end in disaster. In order to learn about the problems generated by "all-up" processes, NASA introduced procedures to document all inflight anomalies in Mission Evaluation Reports (MERs). Over time the numbers of identified anomalies steadily grew into the thousands. NASA policy required that before each new flight, a Flight Readiness Review (FRR) had to confirm that all of the anomalies previously identified in MERs had been resolved.

The overall result was a "can-do" attitude at NASA that enabled it to very rapidly develop many new technologies. Vaughan (1996), 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" (Vaughan, 1996: 209).

After the success of the Apollo program, however, the glamorous vision of putting a man on the moon and securing space flight superiority over the Soviet Union had been accomplished and so was past and it became increasingly difficult for NASA to obtain resources from Congress (The CAIB Report, 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.”

> 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. (Vaughan, 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 reviews] 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 school children 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.

---

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

To some extent, however, a task-performance organizing mode emphasizing predictability masks this complexity because rather than using distributed knowledge to acknowledge and then explore anomalies, procedures are often designed to 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 despite the known presence of thousands of hazards, 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 (The CAIB Report, 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 promises implied a high ability to maintain predictable task-performance – similar to what one might achieve, for example, in managing assembly line production.

In sum, NASA operated in two modes. As Ron Dittemore, the Space Shuttle project manager testified on March 6, 2003:

---

[2] Of the 4,222 are termed "Criticality 1/1R", 3,233 have waivers. CRIT 1/1R component failures are defined as those that will result in loss of the Orbiter and crew. 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.

> "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." (The CAIB Report, 2003:20)

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 1 charts how for the first 9 days of the 14 day flight, events unfolded and interactions occurred between different NASA groups, various technologies, procedures, metrics and other tools that comprised the elements of NASA's distributed mission control and response system. Figure 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 response system.

-- Figure 1 here --

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 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. 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 thermal protection system (TPS). Over time NASA has learned that by making repairs after each flight, it can take care of this damage done to the TPS.

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 they were 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 before and they

---

[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

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," 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.[4]

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 then use to respond to this and other emerging events. The CAIB Report (2003:121-174) suggests that from the beginning, those in mission control interpreted the foam loss event as being in the "in-family" category, a "turnaround" rather than a "safety-in-flight" issue that could be dealt with after the flight had been successfully completed. They were often most concerned about how actions and delays to deal with the issues apparently arising from flight STS-107 might reduce NASA's ability to meet downstream schedules and commitments.

---

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.

The CAIB Report also reports 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 preprogrammed 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 classified as 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, the team would have less authority to ask questions or get things done, and their report could be scheduled to be made on a specific but less urgent date.

In the case of STS-107, some groups involved in NASA’s distributed response system classified the falling foam as an out-of-family event while others including Mission Control classified it 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 – that 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. Nevertheless, 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." (The CAIB Report, 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 to DAT. 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 done by DAT was that the debris might well have compromised the underside of the wing generating an "out-of-family" event that on reentry 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 standards for assessing the analysis. These steps led them 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” (The CAIB Report, 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 (The CAIB Report, 2003: 153).

DAT did not realize that in canceling the request, Ham had not known that it was their 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 effectively 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." (The CAIB Report, 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 e-mail 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." (The CAIB Report, 2003: 157)

Despite the fact that parts of the distributed response system was operating to find out more about the damage the foam strike had caused, the overall system was prematurely stopping exploration efforts and tipping towards 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." (The CAIB Report, 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 identifies 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 does not slip through the system and spin the community up about potential problems that have not been fully vetted through the proper channels." (The CAIB Report, 2003:159)

The Response On Day 9, DAT made its formal presentation to Mission Control. A standing-room-only audience of NASA engineers 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, as was required given NASA's established traditions. The Mission Control Team then saw no reason to change their view held all along that the foam strike was not a safety-of-flight issue. According to the CAIB report (2003:160), 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."

Epilogue

After a series of interactions between the different parts of NASA's distributed response 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 reentered the earth's atmosphere on February 1, hot plasma breached the RCC tiles and the shuttle disintegrated. The CAIB report said:

> "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." (The CAIB Report, 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 response system could have tipped either way. How did the system tip towards a predictable task performance mode that prematurely closed 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 towards this goal. This commitment almost certainly encouraged organizing in support of a predictable task performance mode rather than an 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 it to counter the implicit lack of support for the exploration efforts that they 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 among these, it was certainly 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 any of 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

Organizations are responsive systems – in real time, situations 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 and lead to responses that are appropriate to this perspective.

Timing and temporality are vital in this unfolding 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. This is because doubts and ambiguities about emerging events are inevitable in a context where alternative operating modes are contending to determine the perspective that will determine organizational responses. The resolution of doubts and ambiguities depends not only on which operating mode is raised, but also upon when this perspective is raised. As a particular perspective becomes more prominence, as it is supported by more internal organizational power, and as it has more direct contact with emerging phenomena, so its representatives are in a stronger position to advocate their perspective as being the one that is most appropriate for classifying and dealing with the emerging event.

It is often difficult if not impossible for the people in an organization to determine the status of an event objectively in real-time. This is because organizations are often operating in multiple modes. Events can be understood and acted upon only given a particular perspective. To the extent that an organization is operating in a dual mode, critical ambiguities emerge reflecting these alternative modes that prevent an organization from generating shared "collateral experiences" (March, Sproull and Tamuz, 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 could not be resolved within a short time frame. Despite access to vast resources and widespread goodwill, NASA could not in real-time negotiate an agreement between its alternative operating modes that identified the significance of the foam loss event. The resulting data indeterminacy is one explanation for NASA's inability to act in real time.

These observations and processes have implications for learning in organizations. In normal 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 operating modes, 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 its developing technologies.[5] As NASA is operating in two modes, simultaneously, processes generating indeterminacies may be inherent in the "One NASA" vision.

CONCLUSION

Our study raises disturbing implications for theories on decision-making. The mainstream perspective on decision making is based on the premise that individuals are omnipotent information processors – in real time, they can calculate probabilities and risks and eventually figure out the best course of action to pursue. Faced with ambiguous data or conditions, they may copy one another. Whether through calculation or through emulation, this perspective holds that individuals will somehow eventually converge onto similar understandings leading to similar decision paths.

Our study challenges these mainstream observations in several ways. First, a distributed approach to response anthropomorphizes decision making by bringing in metrics, routines and technologies. Second, rather than triggering a convergence of views and decisions, ambiguous situations can trigger divergences of opinions that can amplify over time, confirming the heterogeneity that characterizes organizations. As a result, data that is generated in real time can remain indeterminate.

---

[5] In their analysis of scientific and technological challenges associated with the construction of aircrafts capable of attaining satellite speeds, Gibbens et al. (1994: 20-21) point out that "discovery in the context of application" generates fundamental discontinuities with previous experiences.

REFERENCES

Beunza, D and Stark, D. 2004. *Tools of the trade: the socio-technology of arbitrage in a Wall Street trading room* Industrial and Corporate Change 13: 369-400

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

Brunsson, N., 1985. *The Irrational Organization,* Chichester: John Wiley & Sons.

Callon, M. 1998. The Embeddedness of Economic Markets in Economics. M. Callon, ed. *The Laws of the Markets*. Blackwell Publishers, Oxford, UK, 1-57.

Garud, R. and Rappa, M. 1994. "A Socio-cognitive Model of Technology Evolution" *Organization Science,* Vol. 5, No. 3, pp. 344-362.

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

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

Johnson, S. B. 2002. *The Secret of Apollo: Systems Management in American and European Space Programs*, Baltimore: Johns Hopkins Press.

Kahneman, D., & Tversky, A., Subjective probability: A judgment of representativeness. *Cognitive Psychology*, 1972, 3, 430-454.

March, J. G. & Simon, H. 1993. *Organizations, Second Edition,* Oxford: Blackwell

March, J. G., Sproull, L. S. and Tamuz, M. 1991. “Learning From Samples of One or Fewer”, *Organization Science*, Vol. 2, No. 1, 1-13.

Murray, C. and Cox, C. B. 1989. *Apollo: The Race to the Moon*, New York: Simon and Schuster.

Perrow, C., 1999 *Normal Accidents*: Living With High-risk Technology, Princeton, NJ: Princeton University

The CAIB Report. 2003. http://www.caib.us/news/report/volume1/

The CAIB Report. 2003. http://www.caib.us/news/report/volume6/

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. 1993. “Collective Mind In Organizations: Heedful Interrelating On Flight Decks,” *Administrative Science Quarterly*, 38, 3, pp. 357-381.

Vaughn, D. 1996. *The Challenger Launch Decision*, University of Illinois Press, Chicago, IL.

Table 1a: 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

Table 1b: Distribute 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

Table 1c: Distributed Arrangements For Task-performance 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

Figure 1:Partial Response Graph of STS-107 Disaster

	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 reentry	
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 reentry	
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 DAT 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		Because it is not designated a “Tiger team”, DAT ‘issues are not given priority by mission control managers

									certain it was a safety-of-flight issue.		
Program Requirements Control Board	Declared bipod foam loss an “action,” not an “inflight 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 a 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 check lists to be explored given out-of-family event		Not designated								Never activated
	Prior to incident	Day 1	Day 2 & 3	Day 4	Day 5	Day 6	Day 7	Day 8	Day 9	Day 17	