A High Reliability Management Perspective on the Deepwater Horizon Spill, Including Research Implications

Emery Roe[i] and Paul R. Schulman[ii]

Abstract

Looking back on events leading up to the Deepwater Horizon Spill raises questions about how the Spill is currently portrayed. We run the risk of focusing too narrowly on a sequence of particular events as the main culprit when the accidents and failures that did not happen are just as important to understand in drawing the right policy implications. *Looking forward* means we also have to better understand how and why "better regulation" can't be the only or even primary answer to more reliable deepwater drilling. Research is proposed detailing these issues from a high reliability management perspective. That research would include identification and tracking of indicators of real-time reliability from rig and drilling operators in the deepwater drilling and exploration sector.

[i] Center for Catastrophic Risk Management, University of California, Berkeley, CA.
[ii] Professor of Government, Mills College, Oakland, CA.

Table of Contents

1. Introduction

We may never know all the events that led up to the Deepwater Horizon Oil Spill, but that will not stop decisionmakers from having to explain what happened. The outlines of their storyline are already clear. The Spill was the culmination of multiple failures, human and technical.[1] What could have been prevented by way of better technology, drilling practice, industry standards, corporate responsibility and government regulation wasn't, and as a result a catastrophe like this became all but inevitable.

The problem with the storyline isn't the facts about the casing, blow-out preventer, cement, saltwater replacing the mud, and more. The real issue for us is the premature policy implications that are now being—and are going to be—drawn from this sequence of mistakes. A blame scenario built around mistake-after-mistake is one that invariably scapegoats the lack of government regulation as the ultimate culprit. For it is government regulation that ensures oversight to prevent error, and a storyline that is a catalogue of preventable error after preventable error is one that nearly always ends up at regulation's front door.

You see this kind of blame scenario everywhere today. Just as it is now argued that mortgage companies and investment firms were "acting rationally" in the face of government's failure to regulate the banking and finance sector, so too increasing numbers will be arguing that oil companies were actually under-regulated and thus "acting rationally" given the incentives and disincentives they faced for deepwater drilling. Just as we are told that the financial crisis requires international coordination and regulation—which of course won't happen any time soon—so too we will hear the argument that oil drilling really requires a powerful national if not international regulatory regime, otherwise we should expect such spills to continue. The soft underbelly of this kind of reasoning is patent: It implies that the original errors were preventable, if and only if—miracle happens here—"effective regulation" were in place at the scale needed from the outset.

More effective regulation, it seems, is required, and certainly its lack had a role in the Spill. But it is too early to claim "bad regulation" as the disease and "good regulation" as the antidote when it comes to catastrophes such as the Deepwater Horizon. The conclusion is premature for at least three sets of reasons.

First, we know that high-level executive management failures in the context of severe market pressures can induce large-company technological disasters. Charles Perrow, the sociologist and complex systems analyst, reviewed the case material and concludes there is "very good evidence of what will be called 'executive failures,' where the chief executive ignores the warnings of his staff and forces them to conduct unsafe practices." [2] But large companies facing the same regulatory incentives and disincentives act in very different ways when it comes to being reliable in what they do. Such differences are not reducible to "regulation." [3]

Second, even with political will, the proffered solution—more effective regulation on a wider scale—isn't the basket in which you would want to put your most important eggs. We know that deepwater oil drilling is incredibly complex and uncertain. By differentiating regulatory responsibilities or creating new or expanded multinational or international regulatory agreements, we instantaneously increase the complexity and uncertainty of regulatory oversight and inspection. Regulatory responsibilities can be confused as well as sharpened by agency differentiation. The larger the territory to be regulated, for example, the more important site-to-site differences are when it comes to interpreting or modifying the regulations.

The remainder of this paper looks to a third set of reasons why it is premature to draw policy implications at this stage from the Spill, and suggests a line of research needed to inform discussions about future regulatory strategies.

Our specific concern centers on a different problem in the "mistake-after-mistake-inevitably-leading-to-disaster" scenario. It gives the impression that the "correct" understanding of the Spill follows this format: a set of identifiable preventable mistakes is discovered, which because they weren't prevented, meant that causally, disaster was inescapable. This conclusion may be part of a carefully researched conclusion, but that conclusion would also have to consider:

- That other mistakes may have been prevented by control operators and support staff (on the rig in the driller's chair or in Houston[iii, 4]) which could have led to an earlier disaster, had they not been prevented;
- That an earlier or later disaster could have happened even if the sequence of events leading up to the Spill had not occurred; that is, the highlighted mistakes were sufficient but not necessary; and/or
- That there was a point at which the control operators and support staff, on the rig in the driller's chair or in Houston, if listened to, could have prevented the blowout from occurring—albeit they may have been unable to prevent another failure from occurring in its place.

The three possibilities can render conclusions about the preventability of the mistakes and/or the inevitability of the Spill much less clear-cut than supposed. In particular, we suspect one reason why regulators always play catch-up to changing events ("Why do we keep fighting the last war?") is that political and regulatory decisionmakers seize on one sequence of mistakes as the "cause" of the catastrophe without having to think through all the other possibilities that were at work at the same time.

Below we argue that you first have to understand what control room operators do and are expected to do, an approach we have examined is High Reliability Management (HRM). The rest of our paper parses the Spill through the HRM approach focused on control operators and support staff, so as to draw out the above possibilities more clearly. Once drawn, we examine their policy implications and compare them to those currently discussed.

[iii] "Just like NASA, BP has a control room in Houston where engineers and managers stare at monitors showing remote craft in a distant location attempting to pull off extraordinary feats." (Gapper)

2. Control Rooms and High Reliability Management

To be clear, the little we know about control operators on the Deepwater Horizon drilling rig and support staff in Houston is what we have taken from media accounts, email interchanges with DHSG members, and published documents related to the spill.

Certainly, there is no single large control room on the rig in the same way as there is in a major nuclear power plant or electricity transmission center. The rig's driller has controls and video screens and some support staff—for example, the mud engineer at another console—to help. There is a separately located control room on the rig to keep the rig afloat, on site, and undertake other operational activities, although this is the domain of the captain and not the driller.[iv] As for Houston, its support operations have been much more in the form of an improvised incident command center than as a permanent control room with shift operators at formal consoles. The organizational and business tensions between Houston and the rig, and within the rig's corporate and drilling components, remain to be more fully accounted for and explained.

Our own professional knowledge about control operators comes from what we have found in our research on water and electric transmission control rooms as well as what others have written about operators in transportation (air and rail), telecommunications and other critical infrastructure control rooms.[v, 5, 6, 7, 8, 9, 10, 11] What follows is based on that knowledge and we would expect that what follows would have to be modified in some respects for drilling control operations, if only because of their off-shore/on-shore components.

Control operators are doing their job when their unit manages the safe and continuous provision of a critical service deemed necessary in a wider organizational or social context. In the context of the oil and natural gas sector, exploration and drilling are considered essential activities. To do that job requires situational awareness of those in the drilling control units.

Why is situational awareness important? Because operators see the technical systems they manage as full of major accidents waiting to happen, which would happen if they were not prevented by the operators with the tools they have. For our purposes, this capacity for high reliability management has four components: (a) key skills of control operators, (b) their avoidance of an accident precursor zone, (c) access to multiple performance modes, and (d) bandwidth management.

The *key operator skills* are the operators' abilities, on one hand to recognize patterns (relationships and trends) at the system level, while on the other hand, formulate specific contingency scenarios for what they are doing right now, right here. Without these skills and the knowledge that comes with knowing system patterns and local contingency scenarios, the control room can't tell what's abnormal or unusual—and that is where the precursor zone comes in. "If we've learned anything so far about the deepwater Gulf of Mexico, it is that it contains surprises. And that means an operator

---

[iv] Peter Marshall, personal communication; William Gale, personal communication.

[v] The literature on operators working in control rooms or in other units operating under mandates for real-time reliability is wide but dispersed. Much has done in the field of Human Factors Analysis and Quality Control, though by no means do these literatures speak with one voice (Roe and Schulman). The reader seeking to delve more in this area by way of material that joins concepts and case material can start with the sampler of Kahneman and Klein, Klein, De Bruijne, Woods and Hollnagel, Garbis and Artman, and Roth et al.

needs depth — depth in terms of resources and expertise — to create the capability to respond to the unexpected," wrote the then BP vice-president for Gulf deepwater developments.[12]

The control operators' *precursor zone* is a set of conditions that are at the limit of the knowledge and skills the operators have with respect to the system they are operating. This edge of control operations is where the skills of operators and their support staff are significantly challenged in relation to their task requirements. Operators are no longer clear on what actions of theirs could lead to accidents and failure versus what would lead to them.

Of course, given time and support staff, operators try to figure out what is going here, but all too often that those assets aren't available. As such, operators avoid this zone as much as they can, not just because the risks are higher from being there but for the more basic reason that they can't actually calculate and compare operational risks once they're there. They end up outside their domain of competence and that is a risk all of them are trained to avoid.

This isn't to say that operators never have to operate at the limits of their knowledge or that when they operate in their own domain of competence everything is smooth and unproblematic. On the contrary: In our research, we found operators had to have access to *multiple performance modes* if they were to be reliable. They continually face situations with varying degrees of unpredictability and/or uncontrollability (they can predict the weather but can't control it when it goes bad). Furthermore, those situations are faced with varying resources and options with which to respond.

Sometimes the volatility (unpredictability and/or uncontrollability) they face is high or it is low, while the options they have with which to face that volatility can be many or few. Behavior that is "just-in-case" (having lots of options even when volatility is low) doesn't work when operators have to be creative "just-in-time" (that is, when system volatility suddenly increases because of, say, bad weather). To be reliable is to provide safe and continuous services even as (especially when) conditions become more unpredictable or uncontrollable, while options grow fewer.

We found that the worst situation for control operators to be in is a prolonged state of "just-for-now" performance where volatility is high and options are now very few. Here doing something to make one thing better can make another worse off. In these situations, operators are constantly resorting to quick fixes and band-aids to hold things together in contrast to those situations of "just-in-time" performance, where their multiple options and skills enable them to assemble and improvise strategies to ensure reliability when conditions temporarily turn bad.

From what we have read, the Deepwater Horizon drilling control operators (and their Houston counterparts) were having to operate increasingly outside known patterns and scenarios, well into their precursor zone (if not beyond into experimenting with unknown unknowns), and in part due to prolonged conditions of having to perform "just-for-now" with recourse to only last-ditch measures in the days and hours leading up to the Spill.

And last-ditch they were. "Jimmy Harrell, Transocean's top rig manager, [said]: 'Guess that's what we have those pinchers for' as he left the meeting, having reluctantly agreed to follow BP's instructions" on April 20th, where Harrell, Transocean's offshore installation manager, appears to have been "referring to the shear rams on the blow-out preventer, which were designed to cut off the flow of oil and gas in the case of emergency." [13] The driller on the rig's tower is said to have added, "We'll work this out later;" only by that point there was no more later to work it out and no viable shear rams able to do that work.[14]

Finally, control operators we have observed operate within formal and informal *bandwidths* for their real-time operations. These bandwidths, which have been determined or settled on beforehand, tend to be ranges, limits, or tolerances within which operators drive or navigate the system. In order to keep the system reliable, they may have to breach the bandwidth bounds if they are experiencing a situation not encountered before.

However, the longer they breach the bandwidths or the more often they have to do so or the more those "bandwidths" are not based in actual experience and practice, the more unreliable control operations become. The existence of informal bandwidths derived from control operator experience is especially important, since they supplement what exists in formal manuals and procedures[vi, 15]

For example, an expedited process for the Minerals Management Service to approve permit requests looks to be an ideal formal procedure, if viewed on its own. However, when BP's rushed request of April 15th to revise its drilling plan is approved within 10 minutes of submission to MMS, the request falls considerably outside the range of MMS or BP experience. What is striking to us is the degree to which the regulators involved had themselves been caught up in the "just-for-now" behavior on the rig in its last weeks.[vii]

It is against this HRM backdrop and interpretation of the Spill that the following series of questions and implications take on research importance.

3. First-Order Questions

First, we have to understand the rate of system change and complexity in control room operations.

If the preceding section is correct, then a fuller account of the Spill requires a more detailed description of "normal control operations" on the Deepwater Horizon rig. As others in the Deepwater Horizon Study Group have pointed out,[viii] the rig was operated safely for years with some level of competence. We also need to know when control operators themselves felt things were going wrong over the course of the rig's operation, i.e., "We're outside our comfort zone right now." All this would be helped by assembling existing, real-time indicators of when operators were being pushed to their performance edges and into their precursor zone.[ix, 16]

[vi] For the debate between Transocean and BP over who was following what with respect to the rig's formal emergency manual on April 20th, see Gold and Chazan.

[vii] Of course, the classic charge is that regulators and industry are always in a cozy relationship, so we should expect such "expediting" to the mutual benefit of each. That may be true, but *in this case* the charge may well miss the important point. The hurried approval of the April 15th permit scarcely did any favors for BP or for the Service when it came to ensuring drilling reliability or their own organizational survival. In HRM terms, the regulators and BP were both caught up in a string of just-for-now quick fixes, band-aids, and last-ditch interventions that took on a life of their own.

[viii] e.g., Peter Marshall, personal communication.

[ix] For one example of indicators focusing on reliability violations, see Roe and Schulman.

To start we would need to know what were the specific real-time reliability requirements in the drilling and rig control units and how were they monitored.

Clearly one requirement was that downward pressure exerted by the mud and other means from the top of the well pipe had to balance the upward pressure exerted by gas and other elements from its bottom. "You don't want to push too hard, but you don't want to push too easy. There's—a delicate balance that has to be maintained at all times. . . . The bottom line is. . . it's all about pressure in a well. That's all it's about. You have to balance the pressures between the seabed and the bottom of the well. And it's not unusual to have problems down below with pressure. That's why have all of the sophisticated instrumentation?" reports an informed crewmember on the Deepwater Horizon.[17] But just what did that instrumentation show and in what form and for how long?

The indicators would have to be more than those events described in media reports by way of confrontations and short-cuts driving operations during the rig's last days. We are talking about a great deal more than a record of what happened in the 50 minutes prior to the blowout when there were overt signs something was wrong.

We would need to know much more about the crew rotations (how long on the rig, how long off), since first days back for control operators can pose considerable skill issues. So too we would need to know more about the on-rig shifts (e.g., Did different shifts have different personalities? What was the turnover or downsizing on any given shift? What was the mood?[x, 18] etc). We would need to know more about the support staff for control operators (e.g., Were off-site Houston engineers supporting on-rig operators or directing them? Where support staff contracted in as vendors, and, if so, with what effect? Etc.). What was being monitored and what should have been monitored, but wasn't? Finally, we would have to have a much better idea of what the operators themselves saw as their major performance edges, namely, those tasks or activities which challenged them regularly or periodically.

With a more detailed control operations description in hand, we'd be positioned to illustrate the relevance, if any, of the three possibilities mentioned earlier. If the operators concerned were actually having to operate in their precursor zone or beyond for prolonged periods of time, that means that a major disaster could have happened at any time during that period. For all we know, a major failure could have happened with an earlier loss of well control and not just the one on the day of the Spill.[xi, 19, 20]

Remember, the risk of operating in the precursor zone is that those managing there do not know or understand all the cause and effect that matter. What works by way of helping matters within their domain of competence could end up making matters altogether worse when undertaken in the precursor zone. That the worse didn't happen with an earlier loss of well control could scarcely have been reassuring to the operators: Failing to fail is not what control rooms mean by reliability.

---

[x] According to one crew member: "Towards the end of the well, right before the explosion, [the mood] was getting worse. We knew we had about two more of these small—supposed 21-day wells to do. And here we are on—six weeks later on a 21-day well. But the biggest—complaint was we were gonna have to go back to where we drilled that 35,000-foot well. And that was an issue, because it was such a long helicopter ride."

[xi] On an earlier loss of well control see Urbina. For the loss of well control on the day of the blowout see the interview transcript with Michael Williams, chief electronics technician on the rig.

Moreover—and this point is major—control and support operators on the rig and in Houston may have prevented other major accidents from happening even though the Spill did occur. Operators we observed are sometimes in the position of having to incur a formal reliability violation in order to avoid an even worse reliability lapse; they are forced to make "errors" in order to keep the system reliable. They do this because they are skilled in taking risks to reduce risks. When that works, that is great, and we would need to determine the level of operator and crew competence on the Deepwater Horizon drilling rig in the years and months prior to the Spill to see if such actions were evident.

However, when this risk-taking behavior is transferred from their domain of competence where skills work, into their precursor zone where skills are significantly challenged, it can have disastrous consequences. Operators preventing accidents waiting to happen may find that even worse accidents occur because the earlier ones were prevented in the way they were. Possibly, every day that the Transocean Offshore Installation Manager (OIM) manager and control operators managed the rig without major incident turned out to be one more day that corporate BP representatives felt justified in undertaking even more hazardous activities.[xii, 21, 22]

4. Second-Order Questions

Second, we have to understand the rate of system change and complexity confronting decisionmakers outside the control units.

If you argue that the specific mistakes leading up to the April 20th blowout were preventable, then from a HRM perspective you have to be careful in not simply assuming that: (1) preventing those mistakes would have meant no other major incident would have occurred at the rig, or (2) very important mistakes were not being prevented even though the Spill did occur. If you too readily accept (1) and (2) then you risk the tyranny of hindsight that renders the past more linear and inevitable than it actually was. In actuality, there may have been a great deal more contingency, coincidence and luck, good and bad, at work leading up the Spill than hindsight's rapid connecting the dots suggests.

From a control room perspective, it is difficult to see how the major policy implication becomes primarily one of calling for more and better regulation. Such regulation would require not just oversight, but, if we are correct, deep hands-on, continuous site inspection, with recourse to real-time supervisory control to correct observed infractions.

In reality, the President's "we will trust but we will verify" has to be backed up the added qualifier "… and when necessary we will directly manage." But no regulatory cadre exists with such a wide purview and deep knowledge, and those who have it—control room crews and their immediate support staff—can't stop infractions from taking place when senior executives and corporate leaders are determined to trade off system reliability (which may in fact be an essential part of the corporation's own brand!) for short-term windfalls or moving onto their next senior position elsewhere.

[xii] According to one report, "the lack of a major catastrophe for many years, even though a rig explosion occurred as recently as a few months ago in the waters off East Timor" has led some expert to conclude that this "success of safety may have set the basis for what happened" in the Deepwater Horizon Spill (Hoyos, Samuelson).

These regulatory challenges need careful analysis. But a prior need is to at least identify empirical measures of when control operators move to their performance edges and into their precursor zones. That way, we would be able to see that those movements into the precursor zone constitute substantial proof that "threats to reliability are actually happening, right now, right here." With indicators, we'd also have a chance to see whether new regulations themselves move operators away from the precursor zone as hoped or closer to those performance edges as often inadvertently happens.

5. Third-Order Questions

Third, it is important to understand the rates of change and growing complexity in the technical systems themselves that confront decisionmakers both inside and outside the control rooms.

For it is change and complexity, we believe, that threaten to overwhelm both groups and challenge the industry safety record as never before.[23, 24] For we argue that increasingly we do not know what the "system" is we are managing until "it" actually fails. Only then do we understand in greater depth what the system is we were meant to be managing reliably but in hindsight weren't.

To see how different this challenge is contrast it with the dominant view about how reliability fails in our large critical infrastructures. In this view, reliability is eroded little by little until disaster occurs. Industry standards are breached, and lo! nothing happens, and then pressure builds to lower standards further and so on, until we find ourselves taking risks that would have been unthinkable a few years before. This is called "the normalization of deviance,"[25] a term coined to describe events leading up to the 1986 Challenger Accident, where warnings and concerns about the O-rings were ignored with each successive successful flight…until that 25th flight in the program.

If the problem were movement closer and closer to a performance edge without really knowing just how close to failure you really are, then an obvious solution would be to get back to those types of error-intolerant industry standards that worked so well before standards were eroded. Again, the call for more and better regulation is heard.

But there is a less comforting way to see the Challenger Accident and other more recent large technical system failures. It may be that only after the O-rings failed—instead of all the other things "waiting" to go wrong—that NASA managers were able to really comprehend just what kind of system the Challenger shuttle actually was and the limits it had to operate under. Technical systems are so complex today that they confront their managers with the challenge of not just finding useful information in the midst of all the data available ("information overload"), but also having to know just what information they need but may not recognize when they see it ("cognitive undercomprehension").

This dual challenge is a direct function of system complexity that arises through design and redesign and through workaround on workaround to compensate for inadequate design. In this world, not only do regulators not know what it is that is failing until failure happens, they can't know what the effective resources are with which to respond until that failure makes the real system apparent. We do not know if this dual challenge is underway in the oil and natural gas industry, but we have seen it in the electricity and water infrastructures we have studied.

6. Conclusion and Proposed Research

If we are correct in our argument, then what decisionmakers should be doing right now, right this minute, is not only seeking to contain the Deepwater Horizon Spill; they should also be giving *equal attention and care* to ensuring that all those other operating deepwater drilling rigs are being managed by control operators well within their respective domains of competence and nowhere near their precursor zones. We should be most concerned about instances where units have been and are still performing outside their comfort zones without relief in sight. We must wonder if the reorganized Minerals Management Service can meet this priority challenge.

All of above, however, is hypothesis and speculation. While based on theory and prior research findings, it is only a reasonable argument at best and requires confirmation and modification in light of the particulars of control operations behavior on the Deepwater Horizon drilling rig and its Houston counterparts over the years of the rig's operations and not just for the days and hours before the Spill.

Much of this information trail would have to be reconstructed from interviews with former control operators. Such a description need not all be based on existing operator knowledge, though that too would have to be canvassed. In addition, a database of indicators of real-time reliability, such as changes recorded by instrumentation, would be very helpful. In sum, we would propose:

- Review of the growing documentation on the Deepwater Horizon;
- Interview relevant control operators and managers from the rig and support staff in Houston;
- Identify indicators of real-time reliability from control operators and documentation (be they from the Deepwater Horizon's operation or from other deepwater drilling rigs); and
- Assembly of a multi-year database around these indicators and tracking their movements.

Such information will be difficult to retrieve, if only because of legal considerations, but it is perishable in another important sense as well. We know from research that the long- and short-memories of control operators are themselves "perishable," especially under fast-changing conditions of an industry like oil and natural gas. It is important that key informant interviews, however confidential and anonymous they may have to be, start as soon as possible with operators.

7. References

1. S. Kirchgaessner, "Mystery over crew's reaction times to disaster," Financial Times, May 28, 2010, 3.
2. C. Perrow, The Next Catastrophe: Reducing Our Vulnerabilities to Natural, Industrial, and Terrorist Disasters (Princeton, NJ: Princeton University Press, 2007), 9-10.
3. E. Roe and P.R. Schulman, High Reliability Management (Stanford, CA: Stanford University Press, 2008).
4. J. Gapper, "BP is drilling itself into deep water," Financial Times, May 6, 2010, 11.
5. E. Roe, op. cit.
6. D. Kahneman and G. Klein, "Conditions for intuitive expertise," American Psychologist, 64 (2009): 515-526.
7. G. Klein, Sources of Power: How People Make Decisions (Cambridge MA: MIT Press, 1999).
8. M. de Bruijne, Networked Reliability: Institutional Fragmentation and the Reliability of Service Provision in Critical Infrastructures (Delft, The Netherlands: Delft University of Technology, 2006).
9. D. Woods and E. Hollnagel, Joint Cognitive Systems (Boca Raton: CRC Press, 2006).
10. C. Garbis and H. Artman, "Team situational awareness as communication practice," In A Cognitive Approach to Situation Awareness: Theory and Application, edited by S. Banbury and S. Tremblay (Great Britain: Ashgate Publishing, 2004).
11. E. Roth, J. Multer, and T. Raslear, "Shared situation awareness as a contributor to high reliability performance in railroad operations," Organization Studies 27(7), (2006): 967-987.
12. E. Rosenthal, "Our fix-it faith," New York Times Week in Review (2010): WK 1.
13. Kirchgaessner, op. cit.
14. Michael Williams Testimony, "FUSCG/BOEM Marine Board of Investigation into the Marine Casualty, Explosion, Fire, Pollution, and Sinking of Mobile Offshore Drilling Unit Deepwater Horizon, with Loss of Life in the Gulf of Mexico, 21-22 April 2010," (July 23, 2010), 5-241. See http://www.wadisasternews.com/external/content/document/3043/856507/1/7-23-10.pdf and http://www.c-spanvideo.org/program/294728-1.
15. R. Gold and G. Chazan, "BP tries to shift blame to Transocean," Wall Street Journal, May 22-23, 2010, A4.
16. E. Roe, op. cit.
17. Michael Williams, op. cit.
18. Michael Williams, op. cit.
19. I. Urbina, "Documents show earlier fears about safety of offshore well," New York Times, May 30, 2010; 1, 16.
20. Michael Williams, op. cit.
21. D. Hoyos, "The rig's blow-out preventer holds the key to what went wrong and why," Financial Times, May 7, 2010, 7.
22. R. Samuelson, "Oil spill reveals dangers of success," Washington Post, June 7, 2010, A17.
23. D. Brooks, "Drilling for certainty," New York Times, May 28, 2010, A23.
24. R. Samuelson, op. cit.
25. D. Vaughan, The Challenger Launch Decision: Risky Technology, Culture, and Deviance at NASA (Chicago: University of Chicago Press, 1996).