Title
Error Types, Risk Assessment, & the Technology Delivery System for RESIN

Author
Roe, Emery

Publication Date
2010-01-24

NSF-UCB RESIN

Working Paper Series

Error Types, Risk Assessment, & the Technology Delivery System for RESIN

RESIN Working Paper: January 24, 2010

Emery Roe
Center for Catastrophic Risk Management
University of California, Berkeley

Abstract:

Solving the wrong problem precisely (error of the third kind, E3) means having a very good idea about just what that "problem" is that is being solved in the RESIN initiative. This paper presents a framework for understanding major errors of the first, second and third kinds (E1, E2 and E3, respectively) and their implications for the RESIN Technology Delivery System (TDS) and the Risk Assessment and Management (RAM) methods, namely, the instruments of a Quality Management Assessment System (QMAS) and a System Analysis Risk Assessment System (SYRAS). Two cautions identifying further areas of study conclude the paper.

NSF RESIN Project
Center for Catastrophic Risk Management
Institute for Economics and Business Research
University of California, Berkeley

January 24, 2010

As RESIN develops new RAM methods for interdependent, interconnected, and interactive systems (I$_3$CISs), we can refine our notion of the E3 error involved in current single-infrastructure risk assessment. Our argument is that focusing only on infrastructure-by-infrastructure priorities increases the risk of managing the wrong problem precisely, namely, we miss the challenge of having to manage priorities arising out of interconnected infrastructures. Were we instead to manage chokepoints at the I$_3$CIS level, we would be better positioned to consider reordering intra-infrastructural priorities in light of what are inter-infrastructural priorities.

That said, if we continue to focus on E3 error, we should be able to identify just what are the E1 and E2 errors involved and their implications for I$_3$CIS RAM (namely, the QMAS and SYRAS instruments) and the TDSs that are said to follow from these specific assessments (for more on the instruments, see Bea 2002). Our 2009 interviews, discussions, literature reviews and Sherman Island site visit put us in a position to identify some of the relationships.

II

Start with the classic Taylor-Russell diagram for E1 and E2 errors, in this case applied to the critical infrastructure (CI) for electricity:

		Proper CI Decision	
		Do Not Shed	*Shed*
Actual CI Decision	*Do Not Shed*	Correct Decision Operations & revenues maintained	E1 Error Widespread outage occurs with economic losses
	Shed	E2 Error Unnecessary service disruption and loss of revenue	Correct Decision Possibly cascading outage avoided

Figure 1. Taylor-Russell diagram of E1 and E2 errors within the CI for electricity (Little 2005)

The problem illustrated in the Taylor-Russell diagram is that by trying to reduce E1 (E2) error the CI manager increases the probability of E2 (E1) error. For example, electricity deregulation put the incentive on reducing costs, especially those related to the E2 error of unnecessary service interruptions and loss of revenue, as below:

		Proper CI Decision	
		Do Not Shed	*Shed*
Actual CI Decision	*Do Not Shed*	Correct Decision Operations & revenues maintained	E1 Error Widespread outage occurs with economic losses
	Shed	E2 Error Unnecessary service disruption and loss of revenue	Correct Decision Possibly cascading outage avoided

Figure 2. Minimizing E2 error relative to E1 error within the CI for electricity (Little 2005)

As illustrated, deregulation meant greater pressures not to interrupt service even though that increased the chances of wider service interruptions resulting when load should have actually been shed. Moving the boundary line leftwards (from Figure 1 to Figure 2 placement) means the grid manager increases the chances of E1 error by trying to reduce the chances of E2 error. By seeking to reduce the error of unnecessary blackouts, that manager adds to the chances of widespread outages by not taking the decision to shed load when in fact it is the proper decision.

Note that no *ex ante* optimal trade-off for the CI exists for E1 and E2 error. We can't say if the tradeoff in Figure 1 or Figure 2 is "better." A 50/50 chance of committing either error is not *a priori* better or worse than settling for a greater than 50% chance of committing an E1 error because you want a less than 50% chance of committing an E2 error. The decision to strike the percentage tradeoff one way rather than another is set by the CI's calculations of costs and benefits at any given time.

III

Now imagine a situation where there are multiple CIs, each essentially making its own E1/E2 percentage tradeoffs. Our RESIN interviews suggest that something like this is actually going on and has been for some time.

Imagine also that the critical infrastructures are more or less functionally interconnected. "More or less" means we are talking about both tightly coupled and loosely coupled critical infrastructures, where a failure or accident in one affects the activities of another depending on the degree of coupling between them.

As we saw from the 2009 review of the literature, problems in tightly-coupled CIs can cascade from the one infrastructure into another. This is especially true when the precipitating infrastructures are electricity and telecoms, so integral are these to the operations of all critical infrastructures, including themselves. Our interviews show similar tight coupling when the precipitating infrastructure is a levee that has breached, e.g., the inter-state impact of the Jones Tract breach on regional transportation and water supply in and beyond the Delta.

Not all infrastructures are coupled to each other to the same degree. Some in fact have been designed and/or managed to be loosely coupled so as to prevent these cascades or negative dependencies. Our interviews at the California Independent System Operator (CAISO) underscore the reliability of managing the state's transmission grid has improved through having multiple telecommunication vendors along with an increasingly reliable Internet. The reliability of electricity transmission has as well depended on having multiple generators and transmission lines available should one suddenly become unavailable, say because of a levee breach. For their part, some levee systems have been designed to ensure that one levee breach does not cascade into other levee breaches, except under extreme flooding conditions. In the language of systems theory, looser coupling potentially means the advantage of positive redundancy and subsystem semi-decomposability.

IV

But loose coupling and its advantages take us only so far. First, as just argued, some critical infrastructures are tightly coupled to other infrastructures, whether one likes it or not. Private household generators may one day be as ubiquitious as cell phones and afford the same positive redundancy, but even then it is unclear how those generators could run the nation's cell phone system if required. Second, changes are taking place within formerly loosely coupled infrastructures that are increasing their inter-infrastructural interconnections. Sherman Island interviews reported examples of having to raise power lines to accommodate shipping traffic, closing a deepwater shipping lane because of a levee breach, and of a ship cutting off an underwater gas line. The more shipping traffic there is and the more variable the traffic requirements, the more we can expect such close interconnections. The general point is that the increasing number and differentiation of transactions within each infrastructure can be expected to pose implications for other infrastructures.

Third, and not least for RESIN, the unintended result of each CI building in positive redundancy and subsystem semi-decomposability may well generate an evolving tight coupling among and between the infrastructures. This is what the spatial intersection of critical infrastructures in the Delta pose in the form of chokepoints. Our interviews and site research so far make it clear that each of the systems we are concerned with has historically built its system without real regard to what the others are doing. When building redundant generators, transmission lines and distribution systems is confined to a spatially restricted area such as the Sacramento Delta, where other infrastructures are doing the same—as in siting alternative satellite towers in the Delta or bridges and roads through it—there is a point at which the infrastructures end up co-located next to each other, whether unintentionally or after a point intentionally (“let’s build the compressor near the road where we can access it”).

In short and to bring the analysis back to the earlier concepts, there is good reason to believe that each CI’s benefit/cost calculation as where to strike its E1 and E2 error tradeoff as well as its constant recalculations of that tradeoff in light of its changing technology and market conditions have major implications for the respective calculations of the other infrastructures that are spatially, if not functionally, interconnected.

This conclusion allows us to be more explicit about the E3 error involved in our analysis. Specifically, managing the wrong problem precisely is what happens when E1 and E2 *errors are managed carefully within infrastructures without regard for their inter-infrastructural implications and management requirements.*[1]

V

Now shift the analysis explicitly to this interconnected level of what we have been calling the I_3CIS in order to better see what the non-alignment of benefit/cost calculations across CIs mean for management.

Consider the Taylor-Russell diagram of the E1 and E2 error facing a hypothetical I_3CIS manager of chokepoints arising from a change of CI service (such as induced by a change in the CI's E1/E2 tradeoff), where the CI in question is within the I_3CIS:

		Proper I_3CIS Manager Decision	
		Do Not Change	*Change*
Actual CI Manager Decision	*Do Not Change CI Service at a Chokepoint*	Correct Decision Operations and revenues maintained for each constituent CI at the chokepoint and across I_3CIS	E1 Error Operation and revenue losses incurred within I_3CIS level, e.g., at this chokepoint or elsewhere in the I_3CIS
	Change Service	E2 Error Unnecessary disruption of service and loss of revenue at the chokepoint or elsewhere in the I_3CIS level	Correct Decision Possibly cascading failure avoided within chokepoint and across I_3CIS level

Figure 3. Taylor-Russell diagram of E1 and E2 errors facing I_3CIS managers from CI changes

Within this framework, the best situation is where the actual decision taken was proper whether the CI manager or I_3CIS manager decided, e.g., not to shed load was a the best decision whether taken by the electricity infrastructure or by the chokepoint manager for the I_3CIS in which the electricity infrastructure is found. By the same token, the worst situation is the flipside, i.e., whether it was taken by the CI or the I_3CIS manager, the

---

[1] Another way to think about this is that to focus on E1 and E2 errors is to invoke having to think through deliberately the E3 errors involved. To insist that the E1 and E2 errors, for example, revolve around "shedding or not shedding load" is to conceptualize a certain kind of failure; other kinds of failure exist, however, and have to be thought about before choosing to focus on the decision to shed or not to shed. My thanks to Ian Mitroff for this point.

actual decision proves to have been wrong in that both levels witness an unmanaged interruption of service(s) as a result.

Two complications immediately arise for CI and I$_3$CIS managers, however, when the actual and proper decision for one manager turns out to be the wrong for the other, as in the following figure:

		I$_3$CIS Manager Decision	
		Actual but not proper from I$_3$CIS perspective	*Actual and proper from I$_3$CIS perspective*
CI Manager Decision	*Actual but not proper from CI perspective*	Worst position	What works at I$_3$CIS chokepoint level is bad for CI level; or what is bad at the CI level turns out to be good at the I$_3$CIS level.
	Actual and proper from CI perspective	What works at CI level is bad at I$_3$CIS chokepoint level; or what is bad at the I$_3$CIS chokepoint level turns out to be good at the CI level.	Best position

Figure 4. Modified Taylor-Russell Diagram for CI/ I$_3$CIS interactions

What is “actual and proper” for the CI manager ends up being “actual and improper” for the I$_3$CIS manager, or vice versa. For example, it may be that the actual and proper decision taken at the I$_3$CIS level—reordering pre-existing CI levee improvements within new I$_3$CIS priorities—may not be considered proper at the CI level under its own benefit/cost calculations—namely, what about next year’s priorities for levee system repairs now displaced or delayed by the I$_3$CIS reordering?

VI

Conventionally, the lack of congruence between I$_3$CIS and CI—the non-alignment of individual B/C calculations across CIs—is considered an example of the principal-agent problem. According to this perspective, in calculating the costs and benefits of taking its decision, the CI does not include externalities associated with that decision. Its calculation does not cover the cost (or benefit) to the other infrastructures within the

I_3CIS. As a principal-agent problem, the solution is to insist that the I_3CIS manager internalize these externalities into the price of the services at the chokepoint, since each constituent CI service should reflect its true costs and benefits within the I_3CIS. If, for example, the costs of levee improvements are not shared by the electricity transmission system protected by those levees, then the role of the chokepoint manager is to internalize and apportion those costs among the CIs involved.

But the principal-agent solution raises its own problems. In particular, we cannot assume that management at the I_3CIS level trumps and is always better than management at the constituent CI level. To be explicit, we cannot assert that the RESIN RAM of interest, QMAS++/SYRAS++ at the I_3CIS level, is always more effective than the older QMAS/SYRAS at the CI level. Why? Because there is always a non-zero probability that the I_3CIS manager makes the wrong decision even under the best I_3CIS RAM available.

More specifically, when I_3CIS management fails (i.e., when the I_3CIS manager make mistakes by taking what proves subsequently to be the wrong decision), the fallback has to be that the constituent CIs within the I_3CIS can absorb or bounce back from any shock of such a mistake.

Stay with the levee example. Assume the levee systems priorities for next year are reordered in light of the I_3CIS manager's risk assessment of chokepoint vulnerability within the I_3CIS. Assume there is a subsequent breach in the levee that would have been repaired had not priorities been reordered to another levee in light of the I_3CIS chokepoint analysis. When a levee fails even under the best I_3CIS management, would we not want the electricity transmission system behind that levee to remain reliable, that is, to be as reliable *as if* the transmission system were designed and managed to be loosely coupled and semi-decomposable from the levee system said to be "protecting" it?

In this view, the best safeguard for I_3CIS mistakes is to ensure CI reliability in terms of compatibility, serviceability, durability and safety. The best fallback for the non-zero probability that even the best QMAS++ and SYRAS++ can lead to what in hindsight is the wrong decision is QMAS and SYRAS being undertaken within each CI. The best fallback for the lack a high reliability control room at the I_3CIS level is to have high *reliability control rooms within each of the CIs found at the chokepoint. In other words, a resilient and sustainable I_3CIS* **is** *its set of resilient and sustainable CIs that are able and willing, when needed, to be reliable on their own.*

VII

Now we are positioned to bring all this to bear on the TDS issue. If the preceding paragraph is true—and it seems something like it has to be from the perspective of I_3CIS resilience and sustainability however defined—then the functions of an I_3CIS chokepoint manager are clearer.

First and foremost, the manager would ensure that QMAS++ and SYRAS++ are undertaken at the I_3CIS level. Even then, the I_3CIS manager would have to seek to ensure that QMAS and SYRAS are undertaken in each of the CIs and in ways that better take into account how its tradeoffs affect the other infrastructures that are tightly coupled, not just spatially but functionally, with it.[2] Discussions so far have identified a range of potential I_3CIS mechanisms that could undertake such functions, including joint inspection teams assessing the status of chokepoints in the field, a high-level commission with oversight responsibility for the coordinated operation of the state's chokepoints, and expanded functions of the state's emergency response system.

Other duties of an I_3CIS manager also seem clearer after the 2009 RESIN research. An important factor in institutionalizing any new RAM would be the presence of formal control rooms to undertake such analyses. Here too there are no guarantees, but what we know from the literature on control operators (e.g., Woods and Hollnagel 2006), on high reliability management (Roe and Schulman 2008), and on what Kahneman and Klein (2009) call "high validity environments" (and their importance for the resilient systems we are seeking) all suggest establishing a control room at the I_3CIS level or barring that, ensuring robust control rooms exist at the CI level.

---

[2] While managers within the I_3CIS and within each of the constituent CIs would ideally be undertaking there own versions of QMAS and SYRAS, their respective versions of resilience and sustainability would differ. Within any given CI, resilience is the ability to bounce back and plan the next step ahead with respect to the critical service in question, should the I_3CIS manager, for example, fail; sustainability is maintenance of the critical service over time with or without the I_3CIS manager. At the I_3CIS level, resilience involves in some as-yet-unspecified way undertaking interactive QMAS++ and SYRAS++, even when any given constituent(s) fail, while sustainability is doing so over time.

The “robust” in the preceding sentence is meant to flag the fact that, since many of the nation’s critical infrastructures are privately owned, they are under continuing and insistent market pressures that may well be at odds with ensuring resilient, let alone sustainable, infrastructures. One additional duty of an I_3CIS manager would be to monitor, report and seek remedies when such changes affect the CI tradeoff of E1 and E2 error in ways that have negative I_3CIS implications.

VIII

Let me end with two cautions over what the I_3CIS manager may NOT be able to do with respect to an improved QMAS++ and SYRAS++ and by way of an improved TDS.

1.

Figures 1-4 talk about “proper decisions.” You might think then that a logical duty of the I_3CIS manager would be to specify through the improved RAMs what is a “proper decision” from the I_3CIS level of analysis for EACH of the constituent infrastructures involved in a chokepoint or within the I_3CIS. Such specification could, for example, be in the form of the I_3CIS manager establishing a benchmark standard of operation for each CI or that manager encouraging the respective CIs to set their bandwidths of operation together and in light of each other requirements (thereby leaving to each CI its own adjustments within those jointly set bandwidths).

Some of this is possible and should be encouraged. However, much of this would require such an intensive level of knowledge of CI operations on the part of the I_3CIS manager as to make such I_3CIS interventions inconceivable.[3] If they were feasible, the CIs would be doing some of this coordination on their own already or regulators would have filled

---

[3] This requirement of knowledge-intensiveness also poses a major objection to the principal-agent formulation of the I_3CIS /CI non-alignment of interests. Remember we are talking here of a wide range of infrastructures that are spatially or functionally interconnected, including but not limited to levees, wetlands and other protective eco-infrastructure, water supplies, electricity, transportation, and telecommunications. To think that an I_3CIS manager would ever have the level of intimate knowledge of individual CI operations required to create and secure property rights and markets for “chokepoint services” in ways that identify externalities, apportion their costs (or benefits) across all the CI involved, and insure the internalization of these externalities in the form of prices beggars the imagination. For example, consider just the extreme difficulties in pricing ecosystem services on their own (Palmer and Filoso 2009).

some of the gap by this point. That is precisely what has not happened. Indeed, such a hands-on standard setting, when not micro-managing duty of I_3CIS managers for CI operations runs contrary to the Kahneman and Klein finding—and one supported by our interviews—that CI managers (and by implications their regulators) should be and are reluctant to talk about operations outside the bounds of their infrastructure competence and knowledge.

That said, these are early days in our research and I want to leave the door open with respect to this issue of I_3CIS “better coordination of CIs” above and beyond the I_3CIS management duties just discussed in the earlier paragraphs of this subsection.

2.

My second caution concerns the relationship between E1, E2 and E3 errors and QMAS and SYRAS (with and without the ++).

Without prejudging what the 2010 Sherman Island analysis finds by way of Pf and Cf, assume for argument’s that we find a high probability of levee failure due to storms at Sherman Island, and by implication at other like Delta islands. Even if we found a lower than expected probability, assume the probability of levee failure would still be high, if only by adding in the likelihoods of earthquakes and rising sea level. Assume furthermore that even if it proves difficult to calculate the interactive Pf for a chokepoint failing conditional on the levee failing, the notion of chokepoint reinforces a finding that the Cf of the levee failure would also be quite high when we take into account losses to electricity, telecoms, roads and other infrastructures due to a levee failure.

In sum, say our 2010 analyses find that Pf and Cf are very likely to be very high for Sherman Island levee failure. The logical implication would then seem to be that government should do everything possible to reduce the probability of levee failure and its consequences, or barring that, "Get out of Dodge," i.e., get out of Sherman Island if not the Delta.

The problem, however, is that even were Sherman Island levee failure highly probable and highly consequential, you would still have the percentage tradeoff between E1 and

E2 error. Doing whatever is thought needed to reduce E1 error associated with levee failure—that is, reducing the chances of not taking the necessary interventions when in fact levees would fail without them—means you expand the chances of E2 error—you increase the possibility of spending resources on what may prove to be unnecessary interventions.[4] Doing everything possible to reduce Pf and Cf has to be balanced against the then increased chances of an unknown portion of the targeted resources and monies being spent on unnecessary or even counterproductive measures. This is clearly the case were "Getting out of Dodge" the option chosen, as there is nothing cost-effective about emergencies, and presumably "getting out of the Delta" would have to be seen as a last resort, emergency measure.

Will QMAS++ and SYRAS++ enable I_3CIS decisionmakers not only to estimate better (interactive) Pfs and Cfs, but also to focus their attention on the percent tradeoff involved in strike the E1 and E2 error balance? One could argue, for example, that QMAS and SYRAS (with or without the ++) are inside-RAM, while the E1/E2 tradeoff is ultimately outside-RAM. In this view, the tradeoff is broadly cultural and political in orientation rather than about the measurement of probabilities and consequences of failure, and the risks that follow from such calculations.

True, but only true up to a point—and the preceding subsection takes us beyond that point. Here's the issue.

It's quite clear that there are widely different perceptions over the Pfs and Cf associated with levee failure in the Delta. One finding of the Sherman Island site visit is that some of these differences are genuinely cognitive, not just political. That is, some people—most?—living and/or working on Delta islands do not see the probability of levee failure there as high as, say, the recent Delta Risk Management Strategy (DRMS) report found.

---

[4] One additional reason why the pressure on a I_3CIS manager to reduce E1 error derives from the centrality of the electrical CI to other CIs at a chokepoint. We can imagine a great reluctance on the part of the manager to agree to shedding load if only because of its implications for other infrastructures at that chokepoint and/or functionally interconnected with it.

To see what is at stake, think of the probability and consequences of levee failure as both being high or low, with four possible states,

		Consequences of Failure	
Probability of Failure		*High*	*Low*
	High	1	3
	Low	2	4

Figure 5. Probability and Consequences of I_3CIS Failure: Four Different State Conditions

Assume that a classic RAM analysis, along the lines undertaken in DRMS, leads to a finding that for Sherman Island and the Delta, reality is state condition 1 (high/high). Assume too that people living and/or working on Sherman Island or in the Delta are distributed largely across state conditions 2 and 3. That is, Delta residents and workers either think a major levee failure is a low probability-high cost event (2) or think that levee failures are happening all the time but most always as low-cost events—a wet spot here, a new dip in the levee road there, or something else but almost all the time with relatively low consequences (3).[5] We have interviewed no one so far who argues state condition 4 holds.

---

[5] The differences between outside and local experts can be accounted for by any number of political and cultural reasons. Some are addressed in Roe White Paper yy/10. Let me mention one possible factor that seems especially pertinent here.

During our 2009 Sherman Island site visit, two different orientations were identified with respect to the same stretch of levee. From one perspective, the levee was protecting the waterside from Island inland contamination and other negative impacts (e.g., on adjacent freshwater wetlands) due to levee breaching. On the other side, the levee was protecting Island farming and other inland activities from the water flooding in.

There is something here of an inlined management perspective versus outline management perspective. Think of a square shown on a piece of paper: the object could be derived by outlining a square directly onto the paper or by drawing objects or set of objects, which leave behind a blank space in the middle shaped like a square. If you draw it directly, it is outlined; if you draw other shapes that leave it behind as a residual artifact, it is inlined.

In a similar way, the Reclamation District for Sherman Island does its repair and maintenance on the levee inside of the island leaving the waterside to be managed or worried about by others, including the Department of Water Resources, Army Corps of Engineers, or the Sacramento and Stockton ports. The RecDistrict sees the levee protecting the land from the water, DWR and the ports may see the levee protecting the water (e.g., deepwater shipping channels) from the land. When perspectives differ in this way, estimates of Pf and Cf surely differ as well.

In brief, assume our local experts are in state conditions 2/3 while the outside experts are in 1. In terms of the interactive phase of QMAS and SYRAS (with or without the ++), three options pose themselves immediately:

1. outside experts convince local experts to increase the latter's Pfs and Cfs in line with what the former have found,
2. local experts convince outside experts that the Pfs or Cfs are lower than the latter initially found, or
3. local and outside experts meet somewhere in between their initial estimates.

By this point, however, a four option should be clear:

4. the experts, whether local or outside, are striking different tradeoffs between E1 and E2 error. While outside experts are willing to risk a greater E2 error in order to reduce E1 error, the local experts want to be assured that E2 error (unnecessarily costly, if not counterproductive interventions) are minimized when it comes to reducing levee failure.[6]

In other words, what look to be differences over estimates of Pf and Cf by outside and inside experts may also (or instead) be differences over where they strike their respective E1/E2 tradeoffs.

Finally, a longer paper would have to tease out what we mean by "Pf" and "Cf" in the preceding paragraph, In particular, for RESIN purposes Pf is not single element, but rather has two components: the probability of failure due to intrinsic uncertainties (PfI) and the probability of failure due to extrinsic uncertainties (PfE). The DRMS study, for example, focused, only on PfI, and even then, not all of it; DRMS did not address PfE at all. Thus DRMS had a more attenuated notion of what I have been calling "Pf" than the fuller notion of PfI and PfE adopted in the RESIN project. That said, the orientation toward E1 and E2 errors with respect to how "conservative" one is in calculating intrinsic and extrinsic uncertainties remains pertinent. For example, at the I_3CIS level, there is no reason to believe that a constituent CI is risk averse in the same way as other CIs are or for that matter as the I_3CIS manager is.

[6] A longer paper would also spend considerably more time on the differences between local and outside experts with respect to their respective Pfs and E1/E2 tradeoffs. Note here only that local people are much more willing to trade-off many future lives against saving one life today than are professional economists in their net present value calculations with respect to the effects of, for example, global climate change (e.g., compare findings in Thaler and Sunstein 2003 with assumptions underlying the Stern Review 2007).

In sum, one would hope that QMAS and SYRAS, both in their CI and I_3CIS forms, would be amenable to identifying factors of concern relevant to better understanding the costs and benefits taken on the basis of any Pf and Cf calculation.

References.

Bea, R. (2002). Human and organizational factors in reliability assessment and management of offshore structures. *Risk Analysis* 22(1): 19-35.

Kahneman, D. and G. Klein (2009). Conditions for intuitive expertise: A failure to disagree. *American Psychologist* 64(6): 515-526.

Little. R.G. (2005). Organizational culture and the performance of critical infrastructure: Modeling and simulation in socio-technological systems. *Proceedings of the 38th Hawaii International Conference on System Sciences – 2005*, IEEE.

Palmer, M.A. and S. Filoso (2009). Restoration of Ecosystem Services for Environmental Markets. *Science* 325 (July 31): 575-576.

Roe, E. and P. Schulman (2008). *High Reliability Managemen*t. Stanford University Press: Stanford, CA.

Stern, N. (2007). *The Economics of Climate Change: The Stern Review*. Cambridge University Press: Great Britain.

Thaler, R. and C. Sunstein (2003). Liberal paternalism is not an oxymoron. *Universitiy of Chicago Law Review* 70.

Woods, D. and Hollnagel, E. (2006). *Joint Cognitive Systems: Patterns in Cognitive Systems Engineering*. Boca Raton FL: Taylor & Francis.