1990

Risk and Design

James E. Krier
University of Michigan Law School, jkrier@umich.edu

RISK AND DESIGN

*JAMES E. KRIER**

RISK springs from uncertainty,[1] uncertainty invites error, and, since error can be costly, we would prefer to avoid it (provided, of course, that avoidance is not more costly yet). While there is much in the Noll and Krier article[2] about judgmental error under conditions of risk and uncertainty, there is little about ways to avoid it. So avoidance—more accurately, minimization—of error costs is the topic I want to address very briefly and partially here.

As the discussion in Noll and Krier points out, cognitive psychologists in particular have given us fairly systematic insights into judgmental error, and some commentators have seized on this knowledge as a new kind of "market failure." Law professor Cass Sunstein, for example, uses precisely those words in talking about "the enormous difficulties people face in dealing with low-probability events. People tend to rely on heuristics that lead to systemic errors."[3] Psychologist Paul Slovic and his colleagues make explicit what is probably implicit in Sunstein's remarks—that this newfound source of market failure is a good reason for government intervention in risk markets. In a review article on risk regulation, they summarize the research regarding "difficulties people have in thinking intuitively about risk and uncertainty." The results of that research, they observe, "run counter to the traditional presumptions of knowledge and rationality that underlie economic approaches to decisionmaking

* Earl Warren DeLano Professor of Law, University of Michigan. Thanks to Christine Papajohn for giving direction to some of the ideas discussed in this commentary.

[1] Classically, "risk" referred only to instances where probabilities are known, other instances being regarded as "uncertain." See Frank H. Knight, Risk, Uncertainty, and Profit 19–20, 197–232 (1921). But this usage seems to have passed.

[2] Roger G. Noll & James E. Krier, Some Implications of Cognitive Psychology for Risk Regulation, in this issue.

[3] Cass R. Sunstein, Interpreting the Regulatory State 46 ("market failure"), 47 (unpublished manuscript, Univ. Chicago Law School, December 1988).

[*Journal of Legal Studies*, vol. XIX (June 1990)]
 0047-2530/90/1902-0011$01.50

under risk. In doing so, these results challenge the viability of market mechanisms for managing risk and thereby suggest that institutional regulation of risk is needed."[4]

The irony of this remark is that the research findings in question could (about) as well be cited in support of just the opposite conclusion, because in addition to challenging the viability of the market in dealing with risk they also challenge the viability of the *government*, as we shall see below. On a parity of reasoning, then, the findings could be used to argue against government intervention, rather than the other way around. Is this what I wish to argue?

Not really, in part because there are reasons independent of cognitive shortcomings that support the case for government regulation of risk. One of these is externalities, as where A and B transact about risk and reach a decision satisfactory to themselves but not to bystander C, who is not part of the transactional (market) setting and whose interests A and B therefore tend to ignore. The literature surveyed in Noll and Krier suggests that A and B might "mistransact" from the standpoint of their own interests because they process certain kinds of information in counterproductive ways. But even if A and B were perfect information processors, still they could be expected to "mistransact" from the standpoint of C because C's interests, although material to socially appropriate risk decisions, will be disregarded by A and B.

To deal with this externality problem, we typically resort to some sort of regulatory device imposed from outside the market, some kind of program of active governmental control. But even if the presence of externalities justifies these programs, still the programs might fail because of the very same cognitive difficulties offered by Slovic and company as a source of *market* failure. Rather than using these cognitive difficulties to indict government intervention, however,[5] I want to consider some ways regulatory programs might strive to limit their consequences so as to avoid especially costly errors. My conclusion is that the two superficially promising ways I have been moved to discuss are not really very promis-

[4] Paul Slovic, Baruch Fischhoff, & Sarah Lichtenstein, Regulation of Risk: A Psychological Perspective, in Regulatory Policy and the Social Sciences 241, 242–43 (Roger G. Noll ed. 1985).

[5] Whether one is talking about market failure caused by cognitive difficulties, externalities, or whatever, it never follows that the mere presence of these defects necessarily means that government action is justified because—as mentioned in the text—the government can fail, too, and sometimes for the very reasons that the market might. For the sake of argument, I am taking government intervention as a given here, whether or not it happens to be justified, in order to examine how government programs might be improved in certain respects.

ing after all. Risk, at bottom, is risky business, and for now I don't see easy ways around that simple fact of life.

Ubiquity. As I read the relevant literature, cognitive failings under conditions of risk and uncertainty are more or less ubiquitous—they are *human* failings, and so they accompany humans not only into the market but into the government as well. I say "more or less" in the foregoing because, while I am sure that some humans are less apt to make judgmental errors than others, all humans make errors nevertheless, and some of these errors can prove very costly. In other words, even "experts"—a term for professional risk assessors—are vulnerable to the heuristics and biases that plague good decisions in the realms of risk and uncertainty, though they might be less vulnerable than lay people. Colin Camerer notes this problem with experts in his commentary on Noll and Krier;[6] the cognitive psychology literature mentions the point regularly;[7] other literature elaborates.[8]

Expert judgment, then, might represent an advance (though obviously not a costless one)[9] over lay judgment yet still fall far short of the ideal; contrary to the implicit claim of Slovic and his coauthors, expertise is not the easy path around cognitive difficulties because experts, too, face the difficulties and hence, at times, make costly mistakes. One way to cope with this problem focuses on improvements in the selection and training of personnel, for the sake of better experts, but I leave that topic to others. Another focuses on the *design* of regulatory systems, for the sake of better institutions. Design is my interest here.

I shall consider two variations on the matter of design, each of which acknowledges the ubiquity of errors in judgment and tries to avoid the most costly mistakes. In that connection, imagine some proposed technology T that is thought to be (no one is absolutely certain) beneficial in some respects and risky in others, and suppose that the issue is whether to introduce the technology. Two kinds of mistake are possible: (1) weighing benefits against risks, T is on balance a good thing, but we mistakenly decide to abandon it; (2) weighing benefits against risks, T is on balance a

[6] Colin F. Camerer, Comment on Noll and Krier, "Some Implications of Cognitive Psychology for Risk Regulation," in this issue.

[7] See, for example, Paul Slovic, Perception of Risk, 136 Sci. 280, 281 (1987); Baruch Fischhoff, Managing Risk Perceptions, in Issues in Sci. & Tech., Fall 1985, at 83, 91.

[8] See, for example, J. Scott Armstrong, The Seer-Sucker Theory: The Value of Experts in Forecasting, in Tech. Rev., June/July 1980, at 19; Clayton P. Gillette & James E. Krier, Risk, Courts, and Agencies, U. Pa. L. Rev. (in press).

[9] There are, first of all, the direct costs of training and maintaining a corps of experts; there are, in addition, the indirect political costs occasioned by the delegation of decision-making authority to technocratic elites. For a sampling of the extensive literature on this last category, see Gillette & Krier, *supra* note 8.

bad thing, but we mistakenly decide to approve it. Call the first kind of mistake a *false positive* because we ask whether T is too dangerous, answer *yes,* and abandon the technology, but this is incorrect. Call the second kind of mistake a *false negative* because we ask the same question, answer *no,* and approve the technology, but now this is incorrect.[10]

Assume for the present (more on this later) that we are dealing with a class of T's we believe to be of a kind where false negatives will generate systematically higher error costs than false positives. If we are to err at all, then, we should be inclined in these cases to err by abandoning a good T rather than by approving a bad one. Since we know we will err from time to time, we want to *purposefully bias* ourselves against false negatives and in favor of false positives as the best approach in an unhappy world of limited information. Consider several ways we can do this by design.

The Architecture of Institutions. Economic and regulatory systems can be designed with reference to what has been called their "architecture," meaning essentially the relationship of various decision-making authorities within a system. Sah and Stiglitz, for example, consider the architectures of polyarchies and hierarchies. A *polyarchy* is defined "as a system in which there are several (and possibly competing) decision makers who can undertake projects (or ideas) independently of one another. In contrast, decision-making authority is more concentrated in a *hierarchy* in the sense that only a few individuals (or only one individual) can undertake projects while others provide support in decision making." As the authors suggest, the market is a prototypical polyarchy, whereas a regulatory system is hierarchical.[11]

True to the view I stated above, Sah and Stiglitz note that all individuals make errors in judgment;[12] their far more interesting observation is that institutional architecture can affect the nature of the errors. In a market

[10] The terminology of false negatives and false positives tends to be used rather arbitrarily, and so, too, for the related terminology, "Type-I and Type-II errors." Sah and Stiglitz say: "All individuals make errors of judgment: some projects that get accepted should have been rejected, and some projects are rejected that should have been accepted. Using an analogy from the classical theory of statistical inference, these errors correspond to Type-II and Type-I errors." Raaj Kumar Sah & Joseph Stiglitz, The Architecture of Economic Systems: Hierarchies and Polyarchies, 76 Am. Econ. Rev. 716 (1986). For an apparently contrary usage, see A. Mitchell Polinsky & Steven Shavell, Legal Error, Litigation, and the Incentive to Obey the Law, 5 J. L., Econ., & Org. 99 (1989). We follow the environmental risk literature and call Sah and Stiglitz's Type-II errors "false negatives." See, for example, Talbot Page, A Generic View of Toxic Chemicals and Similar Risks, 7 Ecol. L. Q. 207, 219–20 (1978).

[11] Sah & Stiglitz, *supra* note 10, at 716.

[12] *Id.*

with proprietary rights, for example, if one firm decides to go forth with a *T*, no other firm may, whereas if one firm rejects a *T*, some other firm may still choose to pursue it. In a certain form of regulatory hierarchy, on the other hand, if a lower level of the hierarchy rejects a *T*, it is dead, whereas otherwise review continues, and a higher level may decide to reject the *T*. The total *T*'s selected in a polyarchical market will thus consist of all *T*'s accepted separately by all firms, but the total *T*'s selected in a hierarchical regulatory system will consist only of the far smaller set of *T*'s not rejected by some level of the hierarchy along the way.

Markets, then, have a comparative architectural advantage in approving proposed *T*'s, and regulatory hierarchies a comparative architectural advantage in rejecting them. Markets will approve more *T*'s, good and bad, than will regulatory hierarchies (designed as Sah and Stiglitz stipulate), which will in turn reject more *T*'s, good and bad. As a result, the incidence of false negatives should be relatively higher in the case of markets, and the incidence of false positives relatively higher in the case of governmental regulation. The latter is best at rejecting bad *T*'s, for which advantage is paid the price of good *T*'s mistakenly forgone.[13]

I have assumed for now that the price of forgone good *T*'s is worth paying in the particular class of cases under discussion. Given that, the very simplified account of Sah and Stiglitz's argument sketched above provides explicit guidance regarding risk regulation. First, the account makes a case for regulation through the government as opposed to the market because hierarchies outperform polyarchies in avoiding false negatives. (In fact, given the ubiquity of cognitive failure, the Sah and Stiglitz account seems to make a better case for government regulation than does the account of the cognitive psychologists, so long as we hold to the idea that limiting false negatives is the dominant objective.) Second, the argument from architecture makes a case for a particular kind of government regulation, namely, a kind where approval, but only approval, of a *T* by some bureaucratic level allows further review and possible reversal by a higher level; rejection at any level should be final and unreviewable. Third, it follows that there is an obvious case for *judicial* review of regulatory activity so as to add an additional level to the hierarchical architecture. Again, however, review should be of a particular kind: if an agency rejects a *T* on the merits somewhere in the course of its decision process, that decision should not be subject to scrutiny by the courts. Judicial oversight should be confined to ultimate regulatory approvals.

The Burden of Proof. Another way to limit false negatives by design is through careful assignment of the burden of proof. A commonplace ex-

[13] *Id.* at 718–19.

ample comes from the criminal law and concerns not false negatives but false positives. The government is required to prove the guilt of a criminal defendant beyond a reasonable doubt because this makes false positives, erroneous convictions of the innocent, very unlikely. At the same time, of course, the high burden of proof faced by the government also increases the likelihood of false negatives (erroneous acquittals of the guilty), but the argument is that we accept these as the necessary and worthwhile price of protecting liberty—something to which we attach a considerable premium.

If "beyond a reasonable doubt" is a way to limit false positives, then its inverse is a way to limit false negatives. So Talbot Page has suggested that shifts in the burdens and standards of proof can be used in the control of T's that are regarded as especially risky from the standpoint of health, safety, and the environment. The approach in such cases could be, for example, that a T shall be abandoned where there is "at least a reasonable doubt" about whether the T is, on balance, a good thing.[14]

Bootstrapping. If error costs are generally believed in some particular context to be systematically asymmetrical, if there is wide agreement about the direction of asymmetry, and if an important objective is to limit the more costly kinds of error, then the foregoing discussion suggests several ways (no doubt there are others) to achieve the objective by design. Choices among alternative designs would not necessarily be easy, of course, because people might agree on the fact and direction of asymmetry but disagree about its magnitude, and the magnitude of asymmetry is relevant to details of design. (For example, we might opt for one standard of proof—such as proof beyond a reasonable doubt—or one kind of hierarchy if we think that false negatives are way more costly than false positives, but choose a different hierarchy or a less demanding standard of proof—say clear and convincing evidence—if we think they are only somewhat more costly.)

Even if there is agreement about the direction and magnitude of asymmetry, still there might be arguments over issues of design. The architecture of institutions and the fashioning of burdens of proof raise issues beyond the minimization of error costs in states of risk and uncertainty. Yet rarely, if at all, can designers unpackage institutions and tinker with procedural rules for the sake of limiting error costs without altering them in ways that might be regarded as undesirable from the standpoint of other concerns. There is, for example, ideological debate about markets and

[14] Page, *supra* note 10, at 234, 239. See also Talbot Page, On the Meaning of the Preponderance Test in Judicial Regulation of Chemical Hazard, 46 L. & Contemp. Prob. 267 (Summer 1983).

hierarchies—governmental hierarchies in particular—that proceeds on grounds having nothing to do with the relationship between those institutions and types and rates of error.[15] Similarly, rules about burdens and standards of proof may serve symbolic and ritualistic functions again unrelated to avoiding mistakes.[16] But even acknowledging all these grounds for debate and controversy, a focus on design can inform deliberations and generate alternative solutions to problems of uncertainty—so long, at least, as there is fundamental agreement on the existence, direction, and magnitude of asymmetrical error costs and on the importance of avoiding especially costly mistakes.

That's the rub. While I can suppose that everybody is interested in minimizing the total costs of error (so long as doing so is not itself too costly in other, say ideological, terms), I can hardly suppose that everybody agrees about what kinds of errors cost how much. To the contrary, debates about risk regulation in particular typically arise because there is so much controversy about precisely those questions.

Go back, for instance, to my example regarding the regulation of a class of T's, as to which there was agreement—this was *assumed*—that the important thing was to avoid false negatives.[17] In constructing the example, I had in mind T's that pose risks of a familiar sort: ones where the threatened harms, if they materialize, will be catastrophic and irreversible but where materialization is very unlikely (has a very low probability). I had in mind, in other words, so-called environmental risks like carcinogenic chemicals, toxic wastes, nuclear fuels, and so forth.[18] A common view in the case of these risks is that error-cost minimization suggests we should design around false negatives—precisely because they can prove to be disastrous.[19] Those who hold this view could take a cue from Sah and Stiglitz and argue for the kind of regulatory structure and the kind of judicial review discussed earlier;[20] the objective would be to minimize mistaken approvals of unduly dangerous T's.

[15] Some people, for example, might generally favor the market over regulation through governmental hierarchies, even if higher error costs could attend the former, because they believe that the market nourishes, and government intervention threatens, personal and political freedom. See, for example, Milton Friedman, Capitalism and Freedom (1962). Other people might be inclined against market ordering, even if it were thought to be otherwise advantageous, on the view that market regimes unduly empower the well-to-do.

[16] See, for example, Laurence H. Tribe, Trial by Mathematics: Precision and Ritual in the Legal Process, 84 Harv. L. Rev. 1329 (1971).

[17] See the text after note 10 *supra*.

[18] The terminology of "environmental risk" is suggested by Page, *supra* note 10, at 207.

[19] See, for example, *id.* at 233–41.

[20] See the text after note 13 *supra*.

But obviously, when we cautiously *reject* a possibly very dangerous T, we necessarily (if only implicitly) *approve* some other T^* at the same time, namely the preexisting T^* that the new T would have displaced. This preexisting T^* may in its own way be dangerous, too—say because it has bad side effects, or because it controls a bad disease much less successfully than would the new T, or both. If the old T^* is more dangerous on balance than the new T, then caution proves to be foolhardy, and the perfect becomes the enemy of the better.

A large band of apparently well-trained and well-informed people make just this argument. Their claim is that the technological status quo is often not nearly as safe as the proposed technologies that society too often bans in the name of safety. In our terms, their claim is that efforts to limit false negatives in the case of environmental risks come at much too high a price, that we should be more concerned with avoiding false positives and act accordingly. Peter Huber, for example, categorically asserts that newly developed technologies are on balance less risky than existing ones and that highly centralized, large-scale, complex technologies are superior in the same terms to simpler or decentralized alternatives.[21] High technologies can, of course, have catastrophic effects, but simple technologies can, too, so the false negative/false positive problem is symmetrical. Even a bathtub filled with water can threaten disaster, Huber says, because "in a 'worst conceivable accident' the citizens of the nation might line up in front of the tub and drown themselves, one at a time."[22] (The same point has been made, but much more credibly, with respect to such technologies as recombinant DNA.)[23]

Whether a technologically optimistic view like Huber's is correct, I

[21] Peter Huber, Safety and the Second Best: The Hazards of Public Risk Management in the Courts, 85 Colum. L. Rev. 277, 307 (1985).

[22] *Id.* at 313.

[23] See, for example, Stephen P. Stich, The Recombinant DNA Debate, 7 Phil. & Pub. Aff. 187, 191 (1978): "Just as there is a non-zero probability that unforeseen consequences of recombinant DNA research will lead to disaster, so there is a non-zero probability that unforeseen consequences of *failing* to pursue the research will lead to disaster" (emphasis in original).

Enthusiasts of technological advance can be expected to favor the market as a chief means to control the development of technology because—in the terms of Sah and Stiglitz—polyarchies promote innovation, whereas hierarchical regulation might bring undersirable costs: "good projects get rejected in the process of ensuring that bad projects do not get undertaken." Sah & Stiglitz, *supra* note 10, at 726. Even when technological enthusiasts concede the case for regulation, they argue that reviewing courts overseeing the regulatory process should at least defer to agency decisions that *approve* new technologies. See, for example, Huber, *supra* note 21, at 332–35, and contrast the discussion in the text after note 13 *supra*. For a critique of Huber's views in this and other respects, see Gillette & Krier, *supra* note 8.

cannot say.[24] What I can say with utter confidence is that it is as widely held as the opposite view that modern technologies, notwithstanding their genuine benefits, are frightfully risky and thus best approached in very guarded fashion. Hence we can identify a fundamental disagreement that has to be resolved *prior* to settling on an appropriate architecture for the regulation of risk. The disagreement itself reflects uncertainty and might be resolved erroneously. Yet, so far as I can see, about this kind of error institutional architecture can contribute little, if anything.

Exactly the same problem arises when we turn from architecture to another aspect of design—the burden of proof. Consider Page's suggestion, alluded to earlier, that "when the potential adverse effects of an environmental risk are many times greater than the potential benefits, a proper standard of proof of danger under the expected cost minimization criterion may be that there is only 'at least a reasonable doubt' that the adverse effect will occur, rather than requiring a greater probability, such as 'more likely than not,' that the effect will occur."[25] Grant the premise of this statement, and everyone might agree with it, more or less. The premise, though, is itself something in question because seldom do we know for sure that "the potential adverse effects of an environmental risk are many times greater than the potential benefits." As we saw just above, people are of strongly different minds on this general issue. In consequence, there is once again an issue of fact that must be resolved *prior* to choosing any standard of proof, but this issue of fact (like all issues of fact) must *itself* be resolved by an appeal to some standard of proof. Who then shall carry the burden of proving, and to what degree, that a given T is in the class of T's whose risks may be established by a relaxed burden of proof? We enter an infinite regress.

Conclusion. Design offers a good way to deal with some kinds of risks and uncertainties once other kinds of risks and uncertainties are resolved. So once there is agreement, say, that certain specified kinds of technologies are best approached with a view to limiting false negatives, design tells us much about how to contain the costs of inevitable errors in judgment. Unhappily, though, current debates about risk reveal strikingly different attitudes about the pros and cons of various risk sources. Have

[24] I am, however, skeptical. See James E. Krier & Clayton P. Gillette, The Un-Easy Case for Technological Optimism, 84 Mich. L. Rev. 405 (1985).

[25] Page, *supra* note 10, at 239. Page appears to acknowledge the difficulty I go on to discuss in the text. See *id.* at 235: "For many environmental risks it is difficult even to define a candidate which might be or might become a false negative, much less to design an institutional structure which would take the chance of a false negative into account as well as the chance of a false positive. The search for environmental risk candidates and false negatives remains an underdeveloped art, and each case is different."

the government seek to prosecute a citizen, and all of us see precious liberty threatened. But have the government seek to "prosecute" a chemical, and reactions vary in the extreme. For some people, prosecution shows a cautious concern about danger, while for others it reflects a reckless disregard of safety. The truth of the matter is uncertain, and this uncertainty cannot be resolved by design. Design safeguards can be sensibly introduced only after we know the relative costs of false positives and false negatives. Design provides no guidance on the hotly contested prior question, Which set of errors carries the greatest risks?