

**Social Construction of Knowledge and Safety Regulation Policy**

Lav R. Varshney

Science, Technology, and Public Policy

Massachusetts Institute of Technology

December 9, 2006

It is held that under perfect market mechanisms, there would be need neither for product safety regulations nor for product liability litigation. Consumers would be able to assess the safety of products and make rational decisions based on this assessment. Unfortunately, due to the imperfectness of safety information possessed by consumers, the market is far from ideal. A *laissez faire* governmental policy and a *caveat emptor* consumption strategy are not sufficient to guarantee socially desirable outcomes. The area of pharmaceuticals and medical devices, in particular, has seen a significant rise in both regulation and tort action since the 1960s. As per the Food, Drug, and Cosmetics Act of 1938 and its subsequent amendments, firms must establish the safety and efficacy of their products, obtaining approval from the Food and Drug Administration (FDA), before they are allowed to market them. Such a pre-market, direct command and control regulation structure is one way of assuaging the effects of imperfect information; an alternative might be an informational regulation requiring safety information to be provided to the consumer. Any regulatory mechanism, however, requires knowledge to determine its parameters. There must exist some method to determine how safe a particular pharmaceutical is and there should also exist some mechanism to determine how reliable this safety parameter determination is. The debate over regulatory policy in this area often centers on the problem of knowledge appraisal, appraisal of knowledge which can seemingly be stated in scientific terms. In this paper, we study whether public decision making processes can be improved in pharmaceutical and medical device cases, and what impact the legitimation and delegitimation of knowledge claims have on some putative improvements. We use silicon breast implants as a specific product through which to explicate some of the points in our study.

As a starting point, we review the regulatory mechanisms requiring safety and efficacy for pharmaceuticals and medical devices that are currently in place. For conciseness, we will refer to the entire class of pharmaceuticals and medical devices simply as drugs. As we had stated, the basic structure of the regulatory process is the requirement of pre-approval of a drug from the FDA. Unlike in other types of health or environmental policy, where there is a continuum of possible regulatory requirements and thresholds, the FDA makes a binary decision on whether to grant drug approval. This decision is nominally determined by scientific data on the safety and efficacy of the drug. We delineate the costs and benefits in the economic calculus of the command and control regulation. It seems rather obvious that restrictions on drugs with severe side effects that are more harmful than helpful to patients should not be allowed. In other

cases, there is more to say. The main benefit for a consumer associated with the approval of a drug is the availability of a drug that has therapeutic effects without significant side effects. The benefit for the producer is the ability to profit from sales of the drug. The government approval process also grants the producer protection from tort actions by reducing liability and also limits the entry of competitors. Thus in the Stiglerian sense (Stigler, 1971), the regulations may be economically beneficial to producers. As the producers are a concentrated interest in the Olsonian sense (Olson, 1982), they would most likely be able to acquire regulations in their best interest. The cost of a drug approval process for consumers is the delay in obtaining a possibly life-enhancing drug. The costs to the producers include the costs of testing as well as the lost profits during the time the drug is seeking approval. An ideal regulation decision would be instantaneous, approve all safe and efficacious drugs and reject all unsafe and inefficacious drugs. Such a regulation decision, however, is never possible due to the time that is required in testing and the inherent errors in classification when there is uncertainty in knowledge.

The three parameters that are of concern in command and control regulation of drugs are the chance that a potentially beneficial drug is rejected, the chance that a dangerous drug is approved, and the delay in making the decision (Grabowski and Vernon, 1983; ch. 22, Viscusi et al., 2005). In statistics, the first type of error is called Type I and in detection theory it is called a missed detection; the second type of error is called Type II and a false alarm. The amount of delay required in making the decision is called the average sample number of the test (Wald, 1947). Adopting a technocratic view of knowledge appraisal, it is well known from the Neyman-Pearson lemma, that for a one-time, front-end decision making process, the likelihood ratio test gives the optimal tradeoff between missed detections and false alarms. Each piece of knowledge about a drug is to be converted into a ratio between the likelihood that the drug is safe and the likelihood that it is not. All of the likelihoods from each of the individual pieces of data are multiplied together and compared with a threshold. The likelihood ratio test defines a receiver operating characteristic, which is a curve displaying the optimal tradeoff between missed detections and false alarms for a fixed sample number. Once this curve has been drawn, the scientist's work is done. With this separation between science and policy, the policymaker simply picks a point on the optimal tradeoff curve, based on the relative costs that are assigned to misses and false alarms. If the precautionary principle is invoked (Ashford and Miller, 1998), then a choice with a small number of false alarms and relatively large number of misses will be

made. If the opposite view is held (Sapolsky, 1990), then the opposite choice will be made. The relative costs that are assigned to misses and false alarms are, in the technocratic sense, mathematically equivalent to assigning probabilities on whether a drug is safe and efficacious or not, prior to any testing being performed. Thus the precautionary approach is equivalent to a pessimistic view that new drugs will most likely be harmful. Similarly the opposite is equivalent to an optimistic approach with the view that new drugs will most likely be beneficial.

The separation between science and policy that is suggested by the likelihood ratio test is a very clean delineation that puts scientists in the position of supplying knowledge and policymakers in the position of applying values. Let us now examine whether such a separation principle and implied technocratic regulatory structure is truly possible in practice. The first dilemma with the Neyman-Pearson lemma that we encounter is that all of the testing knowledge must be construable in likelihood ratio form. If the knowledge is generated within the boundaries of science, as demarcated by Jasanoff (1987), rather than in the gray area of trans-science, then it seems possible that a likelihood function can be generated from a test. When the knowledge is of a trans-science nature, then it seems infeasible to generate a likelihood function. Even the likelihood function produced by a scientist, however, is wide open to criticism. In debates over knowledge appraisal, a party that opposes the conclusion of the scientist may challenge his/her credibility and make accusations of bias. Although the FDA regulation setting process is not supposed to have an adversarial element, such an element almost surely creeps in. It is a general principle in the sociology of science, even in fields that are not particularly policy-relevant, that “the degree of certainty of a piece of scientific knowledge is, in itself, available to different perceptions, interpretations and presentations.” (Pinch, 1981) To take an example from a field that has not entered much into American public policy debates,<sup>1</sup> Shannon’s noisy channel coding theorem was readily accepted and celebrated by electrical engineering theorists whereas it was dismissed by some mathematicians saying, “the discussion is suggestive throughout, rather than mathematical, and it is not always clear that the author’s mathematical intentions are honorable” (Varshney, 2004). In areas of science, engineering, and mathematics, where there is greater interaction with public policy than there is for information theory, the interpretation of knowledge claims may become quite combative. Given the tendency for the “International

---

<sup>1</sup> The ideological foundations of information theory were, however, the subject of political and cultural pressures in the Soviet Union (Mindell et al., 2003).

Brotherhood of Information Theorists” (Massey, 1978) to withdraw to “the green island”<sup>2</sup> (Vetterli, 2006), despite there not being much adversarial challenge of knowledge claims, one can extrapolate the results in drug testing where there is a great deal of contention. This withdrawal leads to separate communities of scientists that hold similar beliefs, and provides an adversarial system with home bases of scientists that are willing to support a particular view. Thus we see that even within the category of questions that Jasanoff classifies as ‘science,’ there can be contention and doubts cast upon the validity of knowledge claims. Another problem that arises in ‘scientific’ questions related to drug safety is that there can be long-term side effects that do not manifest themselves in routine testing. Although the structure of causes and effects is much clearer in drug safety policy than in environmental regulation and economic regulation policy (Reiner), the lack of observability of long-term side effects remains a fundamental problem in terms of system nonlinearity and feedback effects.

Although a front-end analysis prior to command and control approval was never applied in the case of silicon breast implants because these medical devices were not considered under the purview of the FDA at the time when their use commenced, one can still see many of the challenges to the technocratic knowledge appraisal method arising. In fact, the case is enlightening especially since a command and control directive of banning silicon breast implants was made *ex post facto*. In 1992, after a flood of cases against silicon implant manufacturers entered the courts, FDA commissioner David Kessler banned these implants but also made “weak assurances that removal of the implants was unnecessary” (Angell, 1996). Kessler was simultaneously applying the Ashford precautionary principle in the Ashford-Sapolsky tradeoff and countering knowledge claims that silicon breast implants are harmful. In fact, at the time the decision was made, “none of the epidemiologic studies ha[d] been able to demonstrate a clear link between breast implants and connective tissue disease or suggestive symptoms. This does not mean that there cannot be a link, just that it is too small to have been detected by the studies that ha[d] been done. Possibly, much larger studies will show some risk.” (Angell, 1996). In essence, if a technocratic method were applied, all of the likelihood ratios would have been either unity or favoring safety. This indeterminacy expressed by Kessler and in related

---

<sup>2</sup> “The green island” refers to the color of the cover of the *IEEE Transactions on Information Theory* as well to the insularity of the community of researchers and the difficulty in gaining membership into this community.


knowledge claims made by some camps of scientists, however, was used in legal proceedings to express doubts. As noted by Angell (1996) with regards to the silicon breast implant case:

The law frames questions in adversarial terms, and lawyers see problems as best resolved by controlled argument. In contrast, the scientific method is (ideally) not adversarial, but cooperative, and scientists usually find answers in the slow accumulation of evidence from many sources. The different ways of thinking are so ingrained that they may be virtually unconscious. For example, a lawyer questioning an epidemiologist in a deposition asked him why he was undertaking a study of breast implants when one had already been done. To the lawyer, a second study clearly implied that there was something wrong with the first. The epidemiologist was initially confused by the line of questioning. When he explained that no single study was conclusive, that all studies yielded tentative answers, that he was looking for consistency among a number of differently designed studies, it was the lawyer's turn to be confused.

The fundamental problem in the FDA regulation process is that it is outwardly designed to be a likelihood ratio test and yet adversarial knowledge appraisal elements seep in. In some sense this is unavoidable since the final arbiter in the American system is a judicial process, where the adversarial process reigns supreme. Therefore, biased knowledge is used in a context where it is assumed that data is unbiased. This mismatch also leads to the legitimacy problem in public policy, as expounded on by Jasanoff (1987). During the formulation of public policy in an adversarial system, participants deconstruct and try to delegitimize the science that is used to support the opposing side's argument. When a policy is implemented, however, there needs to be some sort of re-legitimation procedure that provides a scientific basis for the decision that was made, so as not to appear arbitrary or determined solely by politically powerful interests. The problem of legitimation is particularly acute in a drug safety command and control policy since a binary decision is made on a phenomenon that appears to be well understood. As noted before, the causal structure is thought to be very clear in drug safety knowledge: either the drug is safe or it is not. Moreover unlike environmental policy, where a threshold may be changed from say 100 ppm to 90 ppm without the overall policy being delegitimized in the eyes of the public, a change in a binary policy indicates that the original policy was totally incorrect, thereby causing a significant delegitimation. According to Foster's discussion of the Jasanoff view, legitimized scientific knowledge in support of a public policy is actually all that scientific knowledge provides to policymakers. "The 'American Exceptionalism' of placing so much emphasis on scientific evidence is merely a useful 'cover' for policymakers.... the public justification of

American regulatory decisions [rely] on quantitative representations of risk in a manner that [seem] to make the scientific basis of policy explicit. ... US policy must 'be made in public,' [so] it is not surprising that 'officials may find the appearance of methodological rigor especially appealing' because it suggests 'policy decisions are being made in a rational, nonarbitrary manner. ... [S]ocial and political elites refuse to accept 'expert' knowledge as the basis for their 'own political acknowledgment of risk.'" Perhaps this is the main problem in American regulatory policy.

Having reviewed the current drug safety regulation system, discussed the role that knowledge appraisal and legitimacy play, and also established the costs and benefits of several of the major actors in this process, we are in a position to formulate possible improvements to the process. Our first possible improvement tries to form a technocratic solution that takes the intrinsic adversarial nature of policymaking as well as the role that trans-science knowledge, as opposed to science knowledge, plays in the regulatory process into account. In some sense, we are trying to develop a new separation principle that takes into account the factors that made the previous separation principle infeasible. First, let us define the boundary between science and trans-science in mathematical terms rather than the sociological terms used by Jasanoff. In our technocratic setting, we define a scientific question as one where the knowledge produced by a test can be cast as a likelihood ratio comparing safe and unsafe. A trans-science question is one that is seemingly expressed as a scientific question, but where the knowledge generated cannot be expressed as a likelihood ratio. Since knowledge cannot be expressed as a likelihood ratio, a likelihood ratio test cannot be performed to obtain the optimal trade-off between false alarms and missed detections. The knowledge generated by trans-science, however, is usually some quantitative function of the safety of a drug. Consider performing a threshold test based on trans-scientific data; clearly this will be suboptimal in the tradeoff between the chance of false alarm and chance of missed detection. One might expect that bias introduced to data by an adversarial process would reduce the quality of a technocratic decision based on trans-scientific data, as occurs in decision making with scientific data, however this is not the case. Recent work in detection theory shows that intuition is not correct, and that adversarial bias can improve decision performance (Chen, et al., 2006; Chen, et al., unpublished), see also Figure 1. The decision theoretic theorem of Chen, et al. shows (under some technical conditions) that when the



**Figure 1.** Figure taken from (Chen, et al., unpublished), see there for mathematical details of the several curves. An example of receiver operating characteristics comparing the probability of false alarm-probability of correct detection performance for several threshold tests based on trans-scientific knowledge. Curves to the upper left corner are better than curves to the lower right. If scientific knowledge was available, a technocratic decision process could achieve the blue curve. If a non-adversarial, technocratic system based on trans-science knowledge is used, the olive curve is obtainable. If an adversarial, technocratic system based on trans-science knowledge is used, the green curve is achievable. Note that the adversarial knowledge improves performance in the case of trans-science knowledge. An adversarial, technocratic decision process based on scientific knowledge would be below the blue curve and result in reduced performance.

test being made is not the likelihood ratio test, but some other threshold test, then the tradeoff between performance can be improved by adding a mixture of positive and negative bias to the data. That is to say, the adversarial system is actually beneficial for an ideal technocrat making decisions based on trans-scientific data. This counterintuitive mathematical result suggests a separation principle for regulatory procedure as follows. First, the trans-scientist collects some data related to drug safety; as noted by Reiner, essentially all metrics used to establish

governmental regulation are flawed in some way, and hence trans-scientific under our definition. Second, some adversarial bias is introduced into the data, with one side claiming that the data underreports the phenomenon and rectifying it according to its standards and the opposing side doing the same, but rectifying in the opposite direction. Third, a scientist determines a receiver operating characteristic from the mixture of the two biased sets of data. Finally, the policymaker applies the costs of false alarm and missed detection to pick a point on the receiver operating characteristic to determine the regulatory policy. Thus it seems that we have proposed a regulatory procedure that handles the trans-scientific nature of knowledge in policy-relevant areas as well as the inevitable adversarial biases that enter into the process. The “open and ritualized clash of conflicting opinions” (Jasanoff, 1987) does in fact lead to the emergence of a better truth.

One might wonder how knowledge appraisal issues enter into this new trans-science, adversarial regulation procedure that we have outlined. Would the outcome of this regulation setting procedure need to be legitimated so as not to appear arbitrary? The first thing to note is that the bias introduced by adversarial delegitimation is an integral part of the process, the part that improves performance from the olive curve to the green curve in Figure 1. The delegitimation process that improves performance, however, also acts to undermine the cognitive authority of the trans-scientific establishment, since it shows that biased data is better for decision making than unbiased data in the presence of trans-scientific knowledge. If the mathematical foundations of the stochastic resonance phenomenon developed in (Chen, et al., 2006) were well understood by the polity, then it would seem that the cognitive authority of the trans-scientific establishment could be replaced by the cognitive authority of the regulatory bureaucrat who weighs the biased data to establish the command and control drug policy. This replacement of cognitive authority, however, seems rather unlikely to occur in the near future, since the stochastic resonance phenomenon has only recently been discovered. Establishing such a counterintuitive result in the psyche of the average citizen seems like a tremendous undertaking; even statisticians and decision theorists have trouble understanding the phenomenon (Varshney, 2006). Thus we believe even though the procedure that we have outlined would result in a better tradeoff between false alarms and missed detections, the perceived loss in a cognitive authority backing the public policy would make the procedure untenable. Jasanoff’s arguments on the need for legitimation of knowledge claims come to the

fore, and show that improving drug regulation policy is not a simple matter of a deficiency-inclusive decision theory.

Another possible method of improving public decision making for drug safety is to adopt an adaptive decision making procedure. Several arguments promoting the benefits of adaptive policy have been put forth (Zuckerman):

There are two fundamental reasons why evaluation and adaptation should play a role in public policy. First, as programs age, initially-beneficial efforts may become increasingly disconnected from the policy environment, or even lead to perverse outcomes whereby they contribute to the problems they were designed to ameliorate. Evaluation and adaptation allow programs to conform more closely to changing long-term trends. Second, given the success of governments in reducing risks during the twentieth century, those that remain are more difficult to diagnose, and are subject to greater uncertainty as to their effects and mitigation costs. Laws and regulations are promulgated before complete information regarding the effects, consequences, and costs of alternative policies are clear. Evaluation and adaptation correct for the lack of perfect information *ex ante* regarding the sources of a problem, the costs of mitigation, and the effectiveness of policy options. Yet despite these methods' apparent utility, routine, organized reconsideration of policies is not a standard feature of governmental policy-making.

Up until now, we have been considering only front-end command and control regulation structures. By allowing back-end adaptation, the rules of the game change significantly, however the need for policy legitimation remains. If back-end adaptation is to be successful in any sense, the front-end measures must be such that they encourage the production of new and improved knowledge about safety. With regards to silicon breast implants, Angell comments that "the answer to the [safety] question will come in incremental steps, one study at a time, and represent the accumulated weight of evidence from many sources. This is the way medical research works; evidence is accumulated slowly and the conclusion is inseparable from the evidence." For data to be generated regarding the safety of drugs, however, it seems inevitable that the drugs must be tested. Epidemiological studies require that the agent under study exist in the field. Thus when comparing false alarms and missed detections, one must consider that false alarms provide safety information whereas missed detections do not. In our previous discussions, we have also not mentioned the delay in the regulatory process, as quantified by the average sample number. Clearly delay is also an important issue; adaptive policy making may allow a reduction in the delay of drugs to the market.

The purely technocratic method of performing an adaptive decision making procedure with scientific knowledge, the adaptive analog to the likelihood ratio test in the Neyman-Pearson lemma, involves the sequential probability ratio test in sequential hypothesis testing theory (Wald, 1947; Siegmund, 1985). It can be shown that the average sample number of a sequential hypothesis test is much less than the average sample number of a front-end test that achieves the same chances of false alarm and missed detection (Wald, 1947; Siegmund, 1985), thus the delay is much less. The sequential test works as follows. Instead of deciding between two alternatives: 'safe' and 'unsafe,' a three-way decision with 'safe,' 'unsafe,' and 'insufficient information' alternatives is made. If the insufficient information alternative is decided, then another data value is collected. Thus in the case of drug safety, the safe alternative would correspond to allowance without evaluation; the unsafe alternative would correspond to disallowance; and the insufficient information alternative would correspond to allowance with evaluation. The thresholds that determine the boundaries of the three regions can be chosen separately by the policymaker according to the precautionary principle, the optimistic principle, or some other principle.

When starting out the sequential experiments with significant uncertainty about the safety and efficacy of the drug and the sequential test in the insufficient information region, it would seem there would be a legitimation problem. The test is showing uncertainty, and yet the regulator is allowing the drug in the market. Such a situation cries out for a party in favor of allowance to try to legitimate the uncertainty as demonstrating sufficient safety that no further testing is needed, whereas a party in favor of rejection would try to delegitimize the uncertainty as an indicator that the drug should be rejected. A reasonable response to the imperfect information problem in the insufficient information region would be an informational regulation rather than a command and control regulation. Thus consumers and producers would have the benefit of information, but the sequential test could continue. It seems that such an approach might also be suspect under legitimacy considerations. Severe warnings of uncertainty of drug safety would certainly cause consumer action at least akin to the caution displayed in not buying new models of cars or beta versions of software. A further problem with starting out the sequential test in the insufficient information region is the dead hand problem. Once a policy is set at the front-end, back-end adjustments become rather difficult. Among the several reasons, a primary one is that the process of monitoring and reevaluation itself tends to have a

delegitimizing influence. The dead-hand problem suggests that all drugs would remain in the approved, monitored, with informational warnings state permanently. Such a permanent state in the insufficient information regime is also problematic from a legitimacy standpoint, since it basically leaves things in the permanent state of ‘science’ rather than a binary state of ‘policy.’ Thus even without taking adversarial or trans-science effects into account, there seems to be no way for a sequential hypothesis testing procedure to meet the legitimacy criterion.

A final option would be a two-state adaptive decision policy; however this also runs into legitimation problems. The act of switching between approval and rejection engenders a great deal of uncertainty in the polity and reduces the cognitive authority of the decision making process. A prime example of the delegitimizing effect of switching among binary decisions (and thereby acknowledging that the previous decision was totally incorrect) could be seen in the 2004 U.S. presidential election, where ‘flip-flopping’ was seen as severely delegitimizing. This effect is also seen in the safety regulation arena: “As in the case of saccharin, once banned, it is extremely difficult to resurrect a substance or product” (Foster, 1999). Similarly, even though silicon breast implants have been re-approved by the FDA, there does not seem to be much of a resurrection of their market. Again, the necessity of legitimizing knowledge claims serves to incapacitate any suggested improvement to the regulatory procedure.

In this study, we have identified instances where the current front-end command and control regulatory procedure for ensuring the safety and efficacy of pharmaceuticals and medical devices suffers from institutional failures. We have also suggested alternative public decision making procedures that seemingly mitigate these institutional failures, however the need for establishing legitimacy by appeal to a recognized cognitive authority seems to trump these improvements. The process of delegitimation and legitimation of knowledge in public policy making is the fundamental problem that renders the suggested technocratic solutions untenable. In this sense, improving the process is somewhat superficial, and cannot be discussed independent of the social construction of knowledge.

## **Acknowledgments**

The author thanks Prof. P. K. Varshney (Syracuse University) for providing a draft of (Chen, et al., unpublished).

## References

- M. Angell, *Science on Trial: The Clash of Medical Evidence and the Law in the Breast Implant Case*, London: W. W. Norton & Co., 1996.
- N. A. Ashford and C. S. Miller, “Low-Level Chemical Exposures: A Challenge for Science and Policy,” *Environmental Science and Technology*, vol. 32, no. 21, pp. 508–509, Nov. 1998.
- H. Chen, P. K. Varshney, J. H. Michels, and S. Kay, “Approaching Near Optimal Detection Performance via Stochastic Resonance,” in *Proceedings of the 2006 IEEE International Conference on Acoustics, Speech and Signal Processing*, Toulouse, France, 21–24 May 2006.
- H. Chen, P. K. Varshney, S. M. Kay, and J. H. Michels, “Theory of the Stochastic Resonance Effect in Signal Detection: Part I - Fixed Detectors,” to appear in *IEEE Trans. Signal Processing*, unpublished.
- J. L. Foster, *The Dead Hand of Environmental Regulation*, The Center for Environmental Initiatives and the MIT Consortium on Environmental Challenges, Sept. 1999.
- H. Grabowski and J. Vernon, *The Regulation of Pharmaceuticals: Balancing the Benefits and Risks*, Washington: American Enterprise Institute, 1983.
- S. S. Jasanoff, “Contested Boundaries in Policy-Relevant Science,” *Social Studies of Science*, vol. 17, no. 2, pp. 195–230, May 1987.
- J. L. Massey, “Joint source and channel coding,” in *Communication Systems and Random Process Theory*, J. K. Skwirzynski, Ed. Alphen aan den Rijn, Holland: Sijthoff & Noordhoff, 1978, pp. 279–293.
- D. A. Mindell, J. Segal, and S. Gerovitch, “Cybernetics and Information Theory in the United States, France and the Soviet Union,” in *Science and Ideology: A Comparative History*, M. Walker, Ed. London: Routledge, 2003, pp. 66–95.
- M. Olson, *The Rise and Decline of Nations*, New Haven: Yale University Press, 1982.
- T. J. Pinch, “The Sun-Set: The Presentation of Certainty in Scientific Life,” *Social Studies of Science*, vol. 11, no. 1, pp. 131–158, Feb. 1981.
- H. M. Sapolsky, “The Politics of Risk,” *Daedalus*, vol. 19, no. 4, pp. 83–96, Fall 1990.
- D. Siegmund, *Sequential Analysis: Tests and Confidence Intervals*, New York: Springer-Verlag, 1985.
- G. J. Stigler, “The Theory of Economic Regulation,” *The Bell Journal of Economics and Management Science*, vol. 2, no. 1, pp. 3–21, Spring 1971.
- L. R. Varshney, “Engineering Theory and Mathematics in the Early Development of Information Theory,” in *Proceedings of the 2004 IEEE Conference on the History of Electronics*, Bletchley Park, England, 28–30 June 2004.
- P. K. Varshney, personal communication, Dec. 2006.

M. Vetterli, Comments at the Information Theory and its Applications Inaugural Workshop, La Jolla, California, 6-10 Feb. 2006.

W. K. Viscusi, J. E. Harrington, Jr., and J. M. Vernon, *Economics of Regulation and Antitrust*, Cambridge, MA: The MIT Press, 2005.

A. Wald, *Sequential Analysis*, New York: John Wiley & Sons, Inc., 1947.

B. Zuckerman, *Adaption Cases: Food Safety and Banking*.