6-2015

Regulating Ambiguous Risks: The Less than Rational Regulation of Pharmaceuticals

W. Kip Viscusi

Richard J. Zeckhauser

REFERENCES
Linked references are available on JSTOR for this article:
https://www.jstor.org/stable/10.2307/26458524?seq=1&cid=pdf-
reference#references_tab_contents
You may need to log in to JSTOR to access the linked references.

Regulating Ambiguous Risks: The Less than Rational Regulation of Pharmaceuticals

W. Kip Viscusi and Richard J. Zeckhauser

ABSTRACT

The US Food and Drug Administration (FDA) balances risks and benefits before approving pharmaceuticals. But powerful behavioral biases that lead to the mishandling of uncertainty also influence its approval process. The FDA places inordinate emphasis on errors of commission versus those of omission, a bias that is compounded by its desire to avoid blame should risks eventuate. Despite extensive testing, uncertainties inevitably remain. We often learn about the risks of drugs after they are on the market. And there are off-label uses of drugs, which are not part of the initial testing. The FDA shows a strong aversion to ambiguous risks. This is the opposite of what is desirable. For any given initial expected risk level, optimal risk-taking decisions involving uncertainty in a multiperiod world should prefer ambiguous risks and the potential for learning relative to well-established risks of the same magnitude. Therefore, the FDA should capitalize on option value.

1. INTRODUCTION

The regulation of prescription drugs by the US Food and Drug Administration (FDA) is one of the most well-established areas of government safety regulation. Given that consumers lack the expertise to assess the desirability of drugs, the government preempts their discretion and has the FDA approve drugs for marketing on the basis of their safety risks and health benefits. Despite the rigorous nature of the FDA review process, uncertainties inevitably remain with respect to the risks and efficacy of drugs. The way the agency deals with uncertainties often differs from

the way it should. This article examines the strengths and biases of current regulations and proposes ways that the government should deal with the risks and uncertainties associated with prescription drugs.

We should emphasize, notwithstanding our critique below, that pharmaceutical innovation has had a tremendously beneficial effect on health. New drugs have significantly raised life expectancy (Lichtenberg 2011), accounting for at least two-thirds of the .6-year increase in life expectancy of elderly Americans from 1996 to 2003 (Lichtenberg 2013), and have boosted the quality of life of the elderly on various dimensions (Lichtenberg 2012). Drug innovations have also reduced cancer mortality rates by 8.0 percent from 2000 to 2009 (Lichtenberg 2014a). This progress helps to account for the pressures, described below, to provide even more rapid access to cancer drugs.

2. THE FOOD AND DRUG ADMINISTRATION REVIEW PROCESS

2.1. Review Process Procedures

Prescription drugs pose potential risks. To safeguard the public against risky or ineffective drugs, the FDA must approve a drug before it can be marketed, and the agency regulates the subsequent marketing of the drug.

To be approved, a drug must meet standards of safety and efficacy in its proposed uses. Companies must provide "substantial evidence" of that safety and efficacy, which at the final test stage usually requires at least two adequate and well-controlled phase 3 studies. Such studies involve from several hundred to 3,000 adults who have the specified condition or disease. However, the FDA may specify alternative ways in which the criteria can be met (Federal Food, Drug, and Cosmetic Act, 21 U.S.C. 355[c]–[d]).

The FDA's benchmarks for safety and efficacy change over time and differ for drugs targeted at different ailments. To illustrate, the FDA recently proposed loosening the efficacy standards for Alzheimer's drugs. Researchers would not have to demonstrate improved daily, real-world functioning. Rather, the drugs would be subject to a less demanding efficacy standard of subtle improvements in performance on memory and reasoning tests (Kolata 2013; Kozauer and Katz 2013). The desire to strike a balance between safety and efficacy often conflicts with views that safety is the paramount consideration. Moreover, the prospect of setting a lower efficacy standard for Alzheimer's drugs raises the safety con-

cern that potential side effects might not be offset by sufficient efficacy benefits.[1]

The FDA allows the two review tracks made possible by the 1992 Prescription Drug User Fee Act (PDUFA; Pub. L. No. 102-571, 106 Stat. 4491, codified in scattered sections of 21 U.S.C.), which established timeline goals for different drugs' reviews and charges drug companies fees to augment the FDA's resources (Thaul [2012] provides a review of these procedures). Drugs offering minor improvements over existing marketed therapies now receive a standard review, which has a review goal of about 10 months. Drugs offering major advances in treatment or therapies where none existed before now receive a priority review, which has a review goal of 6 months. In addition, the FDA can offer either fast-track or accelerated approval to expedite drugs targeted at serious diseases or currently unmet medical needs. None of these review tracks involve any formal alteration of the safety and efficacy criteria applied to new drugs, although some concerns have been raised that safety may be compromised by an excessively hasty review process (Moore and Furberg 2012).

2.2. Drug Lag in the United States

The motivation for many of the changes in the review process that the PDUFA implemented was concern about a perceived drug lag in the United States due to the FDA's emphasis on resolving the safety and efficacy issues prior to market approval and the balance the agency struck between these competing concerns. Many drugs were available in Europe before they became available in the United States. The United Kingdom approved the beta-blocker propranolol for treatment of arrhythmias, angina, and hypertension before it was approved in the United States; many other drugs experienced a similar history.[2] Two more recent examples

1. A *New York Times* editorial voiced the following safety-related concerns: "Independent analysts need to look hard at whether the F.D.A. [Food and Drug Administration] should lower the bar for these drugs—or should demand a very high level of proof of safety and effectiveness before exposing still-healthy people to possible harm. Even if drugs are eventually approved under this new approach, it will be imperative to force manufacturers to conduct follow-up studies, as required by law, to see if patients benefit in the long run. This is a task they are often slow to perform once a drug is on the market" (*New York Times* 2013).

2. See Viscusi, Magat, and Scharff (1996) and US General Accounting Office (1980) for summaries of the approval dates in the United Kingdom and the subsequent dates in the United States for beclomethasone dipropionate, sodium valproate, cimetidine, protirelin, somatotropin, phospholipids, danazol, disopyramide phosphate, and propranolol. These studies also review drugs approved in other countries but not in the United States that have

are Eloxatin (oxaliplatin), an advanced colorectal, head, and neck cancer drug, which was approved in the United States in 2002 after being approved by at least 29 other countries (see Trowbridge and Walker 2007), and the veno-occlusive disease drug defibrotide, which has been approved in the European Union but not the United States.

Drug-lag concerns have diminished, as the FDA has notably accelerated review time because of both PDUFA and administrative decisions, such as the loosening of the standards for Alzheimer's drugs. The median time to approval in 1993 was 22 months for standard drugs and 14 months for priority drugs (FDA 2002). By 2003, these review times had been cut roughly in half, to approximately 12 months for standard drugs and 6 months for priority drugs (FDA 2007). Given these expedited reviews, Downing et al. (2012) conclude that on average the United States no longer has a clear-cut lag relative to other countries. In a reflection of this narrowing gap, Olson (2013) finds that the percentage of new US drugs that were first approved in the United States rose from 28 percent pre-PDUFA to 40 percent in 1992–97 to 50 percent after the passage of the 1997 Food and Drug Administration Modernization Act (Pub. L. No. 105-115, 111 Stat. 2296, codified in scattered sections of 21 U.S.C.). Rated in terms of the quantity-weighted fraction of new pharmaceutical products sold in 2009 that were launched after 1990, the United States ranked eighth (Lichtenberg 2014b). Although the US shortfall is modest, failing to be the international leader is potentially important in that Lichtenberg (2014b) finds that drug vintage accounts for 73 percent of the increased life expectancy at birth from 2000 to 2009.

2.3. Off-Label Uses

Although prescription drugs undergo a rigorous review for specific uses that are indicated in the drug labeling, there are also situations in which doctors can prescribe drugs for uses not specified by the labels and not the subject of FDA testing. In these instances there is no FDA determination that the efficacy benefits for that use outweigh the safety risks. Neither safety nor efficacy assessments are directed to the off-label use. Such off-label uses are especially common with anticancer chemotherapy drugs. Approximately half of all usage of such drugs is employed for indications not specified on the label approved by the FDA (American Society of Clinical Oncology 2006). Both the National Cancer Institute and

---

been found to be unsafe; the sedative and morning-sickness drug thalidomide is the best known example.

the FDA have recognized that there are some circumstances in which off-label uses can be accepted medical practice or are part of the standard of care for the illness. Cancer and other life-threatening illnesses are particularly subject to off-label drug uses, since the trade-off between safety and efficacy may differ from that for patients with less threatening ailments. A study of off-label prescribing of 160 drugs by office-based physicians in 2001 (Radley, Finkelstein, and Stafford 2006) found 150 million off-label mentions. Each drug mention was obtained from a survey of US office-based physicians whereby a drug therapy was recorded for a patient encounter either in terms of a new prescription or continuation of a previously ordered drug therapy. These mentions were appropriately weighted to reflect national utilization patterns. Within a particular class of drugs, cardiac therapies and anticonvulsants had the highest rate of off-label mentions, 46 percent for each (Radley, Finkelstein, and Stafford 2006). A quite substantial 73 percent of all off-label mentions had little or no scientific support, much less a revisiting of FDA's safety and efficacy requirements. Contrary to usual FDA policies that shun uncertainty and tilt to emphasize safety over efficacy, the agency mostly looks the other way for off-label uses although, as noted below, the US Department of Justice and state attorneys general do not.

The Medicare Improvements for Patients and Providers Act of 2008 (Pub. L. No. 110-275, 122 Stat. 2494, codified in scattered sections of 42 U.S.C.) not only permits off-label drug uses but also requires that Medicare cover drugs to treat cancer that have been approved by the FDA for other purposes. To qualify for Medicare reimbursement, the drugs must meet certain specified requirements (Fiegl 2011). In particular, the off-label use either must be supported by one of the compendia for Medicare reimbursement (that is, American Hospital Formulary Service Drug Information, the United States Pharmacopeia and the National Formulary, and DRUGDEX) or must be medically accepted on the basis of clinical evidence from peer-reviewed journals or other materials approved by the secretary of Health and Human Services. Medicaid reimbursement of drugs approved for noncancer uses and used for noncancer purposes requires support in a peer-reviewed journal but does not require compendia approval (Center for Medicare Advocacy 2010). However, unlike cancer drugs, the medical acceptance of noncancer drugs in peer-reviewed journals does not ensure Medicare reimbursement. Companies are permitted to respond to unsolicited requests for information and in response to such requests are allowed to distribute unedited reprints of studies of

off-label drugs from peer-reviewed journals provided that the company provided no funding for the studies. The FDA guidelines require that, in responding to requests about a drug, a company should provide truthful and balanced scientific information (FDA 2011) and should not include any misleading elements (*United States v. Caronia,* 703 F.3d 149 [2d Cir. 2012]; *Washington Legal Foundation v. Henney,* 202 F.3d 331 [DC Cir. 2000]).[3]

The penalties associated with off-label uses of drugs have generated substantial costs to pharmaceutical companies. Pfizer was fined $430 million in 2004 for defrauding Medicaid on the basis of Warner-Lambert's (which Pfizer acquired in 2000) marketing of Neurontin for a diverse set of ailments such as headaches, bipolar disorder, and alcohol detoxification (Schmit 2004). AstraZeneca reached a $520 million settlement in 2010 for its off-label promotion of its schizophrenia drug Seroquel (Arnold 2010). Novartis settled a series of off-label marketing cases for $422.5 million in 2010 for illegal off-label marketing and illegal kickbacks to health care professionals for items such as entertainment, travel expenses, and speaker programs (Wilson 2010). Amgen paid $762 million in 2012 for off-label use of its anemia drug Aranesp for cancer patients not undergoing chemotherapy (Pollack and Secret 2012). A suit by 36 attorneys general against Janssen Pharmaceuticals (whose parent company is Johnson & Johnson) led to a $181 million settlement for deceptive marketing (Lammi 2012). Off-label uses of the recombinant factor VIIa (rFVIIa) for major hemorrhage have caused concern because of the absence of adequate safety data. Indeed, some studies indicate that rFVIIa does not significantly reduce mortality but does increase the rate of thromboembolic events. Fen/phen, an off-label combination of fenfluramine and phentermine to treat obesity that caused pulmonary hypertension and heart valve problems, spawned considerable litigation and led to billions of dollars in damage awards. It was removed from the market in 1997 (Miller 2009). Such salient negative examples weigh against the benefits for patients and clinical innovations that come from off-label uses (Lipworth et al. 2012; Logan, Yank, and Stafford 2011; Yank et al. 2011).

3. The *Caronia* decision in the Second Circuit gives companies more leeway in providing information.

2.4. Black-Box Warnings, Postlabeling Changes, and Product Recalls

Even approved drugs may pose potential risks. Some of these risks are known at the time of approval; other risks become manifest only later. Drugs posing risks of serious or life-threatening adverse events receive black-box warnings in which the risk warning on the patient packet insert is surrounded by a black box to draw attention to it. A review by Begosh et al. (2006) of new molecular entities submitted to the FDA from May 1981 to February 2006 found that 81 of the 516 such entities received premarket black-box warnings. Many black-box warnings are required after a drug has been approved and marketed (Murphy and Roberts 2006). From 1975 to 1999, 45 of 548 new chemical entities approved by the FDA acquired black-box warnings after approval (Lasser et al. 2002). From 2004 to 2006, the FDA required 1,569 labeling changes, including 174 black-box warning changes—77 were newly imposed black boxes, and 97 were revisions to black-box warnings (Cook, Gurugubelli, and Bero 2009). The remaining 1,395 changes in drug labeling did not involve black-box warnings.

If postapproval risk appearances are grave enough, drugs may be recalled. From 2004 to 2011, there were 91 class 1 recalls, which are recalls of products that have the greatest likelihood of causing patient harm (Wang, Gagne, and Choudry 2012). For the 548 new chemical entities approved from 1975 to 1999, 16 were withdrawn altogether (Lasser et al. 2002).

Olson's recent study of 381 new chemical entities approved by the FDA from 1990 to 2004 provides evidence of a broad set of drug-related risks and their evolution over time (Olson 2013). At the time of approval, 17 percent of those drugs included a boxed warning, and 3 percent included a pregnancy contraindication only in the black box. In the 2 years after approval, there were 308 adverse reactions on average per drug according to the FDA's Adverse Event Reporting System. The distribution of these events is skewed, as the largest concentrations of drugs fall into categories with 80 or fewer serious adverse drug reactions. Drugs on average required 183 hospitalizations and led to 56 deaths. The drugs that were associated with more serious adverse reactions disproportionately included novel drugs and drugs with black-box warnings. Adverse reactions were less common for drugs first launched in the United States, those with a long launch lag after being first launched in another country, and those with a long FDA review time.

3. RATIONAL REGULATION IN A SINGLE-PERIOD WORLD

To simplify our argument, we assume that there are only two outcomes, life and death. The drug offers the prospect of reducing the risk of death but also has the potential to cause the death of someone who takes it.

3.1. Errors of Commission and Omission

Our regulatory policies, and many individuals, regard an error of commission, taking a drug that causes death—as being much more serious than an error of omission—not taking a drug that would have saved one's life. The FDA's emphasis on avoiding type II errors, namely, avoiding the approval of unsafe drugs, rather than type I errors, namely, failing to approve beneficial drugs, has been a long-standing theme in the literature (Viscusi, Magat, and Scharff 1996; US General Accounting Office 1980). The surprisingly modest percentage of FDA-approved drugs that continue to receive their first approvals in the United States and the agency's caution with respect to experimental drugs for gravely ill patients reflect this emphasis (Edwards 2012).[4] One reason for this is that people are subject to reference-point effects and may respond in an asymmetric way to changes in risk, as risk increases receive greater weight. On reflection, however, one might reframe the decision problem: abstracting from cost considerations, should one not follow the strategy that maximizes the probability that one stays alive? In what follows, we frame this question by attaching a weight 1 to being saved by a drug and the weight w to being killed by it. Those who consider errors of commission to be more serious would have a w value greater than 1. An alternate decision calculus, one that we happen to embrace, treats the two errors the same and maximizes the probability of staying alive. It thus sets w equal to 1.

Absent the drug, the individual has a probability s of survival. The drug increases the probability of survival by $i,$ which reflects its therapeutic properties, and reduces the probability of survival by r because of possible adverse effects.[5] In terms of safety and efficacy, safety risks would be $r,$ and efficacy benefits would be $i.$

---

4. The FDA has begun to address these concerns (see FDA, Expanded Access to Investigational Drugs for Treatment Use, 74 Fed. Reg. 40900-01 [August 13, 2009]).

5. There may be instances in which the adverse effect of the drug kills the individual but otherwise would have saved him. In that instance, we learn only that the drug killed him. In most instances, we will never know whether, absent adverse effects, the drug would have saved him.

3.2. The Individual's Risk Calculation

The individual has an overall probability of surviving of $s + i - r$. To simplify the model, we set aside all cost considerations although we recognize that costs are an important consideration in the efficient allocation of health care resources. Posit that he is a von Neumann–Morgenstern (VN-M) decision maker. However, like our current drug regulatory policies, he incorporates considerations of omission or commission in assessing those end states. Arbitrarily let us assign a utility of one to survival and zero to natural death. If reductions r in the probability of survival receive a weight $w > 1$, the expected utility (EU) is thus

$$EU = s + i - rw.$$

An individual for whom $w = 1.5$ would just accept a drug that offers a 15 percent chance of saving him if it has a 10 percent treatment mortality, since

$$s + .15 - .10(1.5) = s.$$

There are three things to note about this formulation in relation to current FDA policy. First, consistent with that policy, it allows errors of commission to be weighted more heavily than errors of omission. Second, it is at variance with current FDA policy, at least as it applies in most areas, in that it employs a consistent trade-off rate between saving and costing lives. The safety requirement of FDA policy in most contexts would create a bias against a drug with a significant probability of a highly consequential negative outcome. This reluctance to approve potentially beneficial drugs that involve high levels of risks is reflected in the continuing limitations on access to experimental drugs for gravely ill patients (see Edwards 2012). This issue received prominence when the family of Abigail Burroughs unsuccessfully sought to secure approval for their daughter to take Erbitux as a treatment for neck and head cancer. The DC Circuit Court upheld the district court's finding that "there is no constitutional right of access to unapproved drugs" (*Abigail Alliance for Better Access to Developmental Drugs v. von Eschenbach,* 495 F.3d 695, 700 [D.C. Cir. 2007], quoting *Abigail Alliance for Better Access to Developmental Drugs v. McClellan,* No. 03-1601, 2004 WL 3777340, 1 [D.D.C. August 30, 2004]).[6] Third, if the FDA were to regulate the drug choices available to patients on an individual basis, the drugs made avail-

6. Also see Stohr (2008) for further discussion of FDA barriers to experimental drugs.

able for people would depend on their preferences. This would clearly lead to significant regulatory complexities because physicians would be told that they could give a drug to a patient but only after carefully inquiring about that patient's preferences. But the decision to proceed with many treatments already often depends on personal preferences.

The role of personal preferences arises with respect not only to probabilistic outcomes but also to unproven safety. Activist groups for sufferers from various diseases—almost always devastating diseases—have sought to influence FDA policy by expressing their passionate desire to have new drugs be made available, even if their safety is unproven. The most dramatic example was for AIDS patients. Two weeks after the AIDS drug protest event Seize Control of the FDA took place on October 11, 1988, the FDA announced an expedited drug approval process for HIV drugs.[7] In other instances, the FDA has in effect channeled the preferences of patients and adjusted its rules to allow drugs that otherwise would have been prohibited. The recent situation with allowing unproven drugs to be administered to sufferers from early symptoms of Alzheimer's disease is a salient case. The insights from common-law approaches—that is, following policies that distill the principles from prior decisions—may infuse FDA rulings more in the future.

3.3. Raisin Benefits versus Watermelon Costs

In some instances, the regulatory process must compare dramatically different benefits and costs, or in effect compare the economic equivalent of raisins and watermelons. Think, for instance, of a hypothetical new drug that enables people to dramatically reduce the pain of arthritis. Suppose that it is also known that one in 100,000 users will have a massive adverse reaction and suffer a fatal heart attack. Should the FDA allow this drug? If not, what if the risk were one in 10 million? Decision theory tells us that there is some probability that would make the risk worthwhile. Indeed, that would be true even if the drug merely partly cleared up acne at the expense of a rare heart attack.

In this spirit, a medical panel recently suggested that women might take tamoxifen or raloxifene to lower their risk of breast cancer even though these drugs raise the risks of blood clots and uterine cancer (Grady 2013). For a group of 1,000 women, the drugs will prevent seven to nine cases of invasive breast cancer over 5 years but lead to blood clots

7. Brier (2009, pp. 165–66) discusses this chronology and the policy changes.

for four to seven women and additional cases of uterine cancer for four women, on average. The desirability of the drugs depends on patients' heterogeneity, including current risk factors such as family history, which would tilt the odds, and preferences relating to side effects and the alternative health outcomes.

3.4. The Food and Drug Administration's Approach to Safety and Efficacy

This discussion of rational regulation roughly captures the spirit of the FDA's practices, but the words "roughly" and "spirit" should be emphasized. Before the FDA will approve a drug, it must pass tests of both safety and efficacy. Exactly how the FDA strikes this balance is not readily observable. The legislative guidance is quite general, and FDA operations are far from transparent.[8] As a result, we focus on our stylized characterization of FDA policy decisions. Let us return to our spare example in which the drug only increases the probability i of survival and independently reduces the probability of survival by r. For simplicity, suppose that the FDA sets standards of safety and efficacy independently. Assume that the FDA would grant approval if and only if two constraints were met: $r < r^*$ and $i > i^*$, where these values may depend on the seriousness of the disease.

The constraints i^* and r^* should relate to each other. The FDA's enabling legislation calls for a balancing of these concerns, but how, and to what extent, there is balancing is not clear. This trade-off was faced 19 years ago, when Sally Zeckhauser (wife of Richard) was discovered to have stage 3 breast cancer. One potential treatment was a bone marrow transplant (BMT). Her doctors said that the procedure has a 4 percent treatment mortality and that they do not like to undertake such procedures. Sally and Richard had consulted the literature and thus asked the doctors for their best estimate of the gain in survival probability from a BMT. Although reluctant to answer, their guesstimate was 10 percent.

---

8. In the 2012 amendment to the Federal Food, Drug, and Cosmetic Act, Senate bill 3187 amended 21 U.S.C. 355(d), adding: "The Secretary shall implement a structured risk-benefit assessment framework in the new drug approval process to facilitate the balanced consideration of benefits and risks, a consistent and systematic approach to the discussion and regulatory decision making, and the communication of the benefits and risks of new drugs." How the balancing should be done is not articulated except to indicate pertinent factors such as "the seriousness of the disease or condition that is to be treated with the drug."
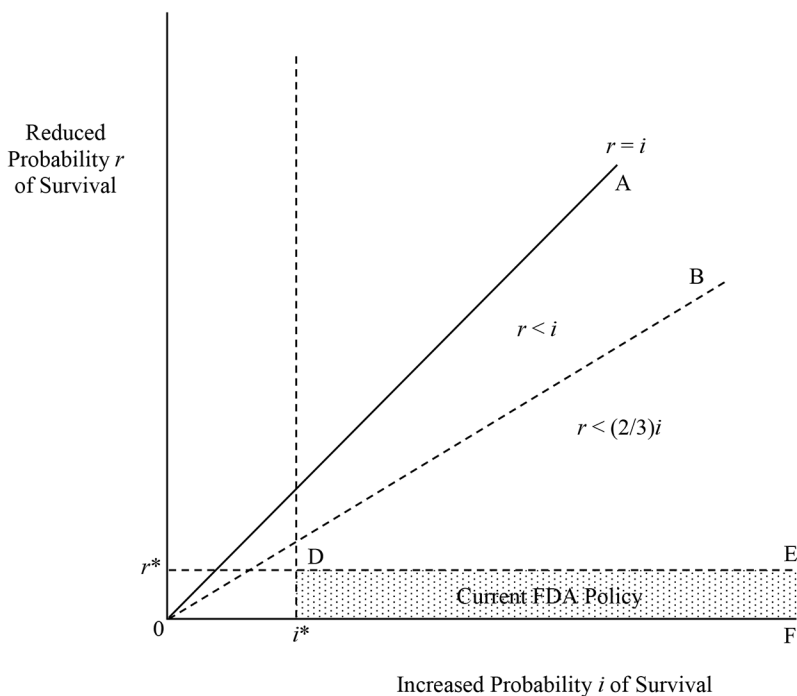
Figure 1. Risk balancing by the Food and Drug Administration

The decision was straightforward. (The 6 percent expected survival gain made it worthwhile going through the agonizing treatment.) Although the treatment pushed Sally to the edge of death, today she is totally healthy.

Our simplified FDA decision framework requires all drugs to be in the shaded area in Figure 1 labeled "current FDA policy." A straight maximization of the probability-of-survival approach would require that the drug be to the right of line A. If errors of commission were valued at 1.5 times those of omission, an approved drug would have to be to the right of line B. The hypothesized zone of current acceptable policies leads to the rejection of many risk-reducing drugs.

3.5. The Food and Drug Administration and a Conservative Approach to Uncertain Probabilities

The losses generated by FDA policies likely are increased by the way the agency treats uncertainty in its establishment of the safety and efficacy

cutoffs.[9] Suppose that the agency does not know the values of r and i with precision but that there is some error around each of them, which we denote e_1 and e_2. If the agency adopts a conservatism principle of focusing on worst-case scenarios, it might view the efficacy or reduced probability of survival due to the drug as being $r + 2e_1$ and the increased probability of survival level as being $i - 2e_2$. Such a conservatism bias has been documented in the case of US Environmental Protection Agency (EPA) policies, in which the extent of the bias varies depending on the worst-case assumptions used. The EPA sometimes uses a series of assumptions in which the upper bound of the 95th percentile is employed for various parameters in the risk calculation. However, when such conservatism is incorporated in numerous parameters that enter the risk calculation multiplicatively, the compounding of the conservatism biases leads to an assessed risk that is beyond the 99th percentile of the actual risk distribution (Hamilton and Viscusi 1999). Although the FDA's conservatism biases with respect to pharmaceuticals are not well documented, the agency has adopted an explicitly conservative approach in assessing the risks of seafood after the BP oil spill (Dickey 2012). Biases of this type may be a more general phenomenon in governmental risk assessment (Nichols and Zeckhauser 1986).

The effect of the conservatism bias depends on which tail of the distribution is used and the manner in which conservatism enters. If the FDA bases its policies on the lower-bound value of both the efficacy and safety levels, that will create a powerful bias against new and uncertain drugs. Such a conservative approach to risk would create a bias against drugs that are potentially beneficial but whose performance is less certain. The worst-case scenarios for drugs require conservatism on both safety and efficacy, say 95th-percentile lower bounds for each. This approach leads to the rejections of drugs that, on an expected-risk basis, offer beneficial net health effects.[10]

The conservatism bias in risk assessment relates closely to the framing of the precautionary principle that is embodied in the entire FDA regulatory approach. In particular, before a new drug can be approved, the

9. In the extreme, the situation may be one of ignorance (Zeckhauser 2010) in which there is no awareness of particular possible states of the world.

10. Moreover, the broader the perceived distributions of benefits and costs, the more costly would be this conservatism. To see this, consider a symmetric extremely tight distribution. Then the 95th percentile is extremely close to the 50th percentile, and virtually nothing is lost by choosing the former rather than the latter. (The 50th percentile, the median and mean, would be the cutoff for a straight maximize-expected-value criterion.)

company must show that the drug is safe. To be truly safe, a drug must be risk free. However, it is never feasible to prove that any risk is 0, as the risk may not have arisen in the sample to date. A demonstration of 0 risk is not the approach of statistical hypothesis testing. An effect of 0 would be the null hypothesis. It is feasible to test the hypothesis that a drug's risk cannot be distinguished statistically from 0, but that is quite different from proving that it imposes 0 risk.

Fortunately, the FDA's standard for safety is not that there be 0 risk. But what level of currently assessed risk the FDA finds too high to approve a drug—that is, the cutoff level—is not clear. How the FDA selects this cutoff level also is not clear.

3.6. The Role of Heterogeneity in the Single-Period Case

Matters become more complicated if the patient population is heterogeneous. To simplify the analysis, assume that all patients suffer from the same disease. Wholly apart from differences in preferences with respect to risk, longevity, and treatment modes, there may be fundamental differences in the treatment-decision problem, such as a disease progressing to different extents in different patients. Thus, patient A might be gaining an expected i_A from the drug, whereas patient B might be gaining an expected i_B, where $i_A \gg i_B$. Posit that A and B have identical preferences. Patient A should accept a much greater risk from side effects of the drug than should patient B.

Stimulated both by concern over the US drug lag and by the spread of AIDS, the FDA altered its regulations to address situations in which patients face life-threatening ailments (Brier 2009, p. 166). However, these policies did not alter the general safety and efficacy tests for the drugs. Instead, the emphasis changed to providing a more expeditious approval process for the drugs, which is quite different from redrawing the admissible drug approval regions shown in Figure 1.

4. RISK AMBIGUITY AND RATIONAL REGULATION IN A MULTIPERIOD WORLD

4.1. Optimal Experimentation with Uncertain Risks

While uncertainty regarding a probability should be a matter of indifference in single-period choices, imprecision in probability assessments of a drug's properties is actually, though perhaps counterintuitively, a desirable feature. The principle that a certain probability is inferior to a com-

pound lottery offering the same expected value applies in a wide range of situations in which a choice can be changed in the future (for example, a drug on the market can be withdrawn or a patient can be put on a drug and then switched to another drug). That is because the compound lottery offers learning opportunities.[11] We illustrate the principle with two simple examples, but the result is quite general.

The first example is stimulated by the classic two-armed bandit problem. Suppose that a patient is faced with two different treatment regimens, which he can select in each of the two periods, with associated probabilities of being healthy or sick. Healthy in each period produces a payoff of 1, and sick produces a payoff of 0. There is no discounting. Treatment 1 offers a known probability of .5 of making the person healthy in either period. Treatment 2 offers an uncertain probability of .5, which we assume is characterized by a uniform beta distribution on [0, 1]. In particular, assume a beta distribution characterized by two parameters, α and β. The person has a prior belief as if having observed α successful outcomes and β unsuccessful outcomes, so the mean perceived risk is $\alpha/(\alpha + \beta)$. With uniform prior beliefs, $\alpha = 1$ and $\beta = 1$, and the mean value of the prior belief is .5. The value of $(\alpha + \beta)$ can be viewed as the measure of uncertainty, with lower values of $(\alpha + \beta)$ indicating greater uncertainty. Which treatment option should the patient pick in each period? The expected value of treatment 1 over the 2 periods is .5 $\times$ 1 + .5 $\times$ 1 = 1. In contrast, learning and adaptive behavior play important roles for treatment 2. The first period of treatment 2 offers the same .5 expected payoff as that of treatment 1. If treatment 2 produces an outcome of healthy in the first period, then the updated probability that treatment 2 produces an outcome of healthy in the second period is .67.[12] Given its success in period 1, the overall expected payoff to this strategy is .5 $\times$ 1 + .67 $\times$ 1 = 1.17. This illustrates the stay-on-the-winner rule for two-armed bandit problems.

However, if treatment 2 produces an outcome of sick in the first period, the probability that it will produce an outcome of healthy in the second period falls to .33. It now becomes better to take advantage of learning and switch to treatment 1, whose probability of producing an

11. Unfortunately, there is a strong behavioral bias that leads such learning opportunities to be neglected (Trautmann and Zeckhauser 2013).

12. With a beta distribution and uniform prior beliefs, the posterior probability of a healthy outcome after a favorable outcome in a single trial is $(1 + 1)/(2 + 1) = .67$.

outcome of healthy stays at .5. Given its failure in the first period, the overall expected payoff to this strategy is $0 + .5 \times 1 = .5$.

The expected payoff for starting with treatment 2 and switching to treatment 1 if the first period produces an outcome of sick is $.5 \times 1.67 + .5 \times .5 = 1.085$, which is 8.5 percent higher than the payoff to starting with treatment 1. The more periods there are, holding other factors equal, the greater is the per-period advantage of the dynamic learning strategy.[13]

In some contexts, individuals need success in multiple periods to have a beneficial result. Think of an individual who will receive 2 periods of treatment, with a life or death outcome in each period. If he survives for 2 periods, he will also enjoy a third period of life for sure. Survival in each period provides a utility of 1, and death provides a utility of 0, as before. If he dies in either period 1 or period 2, he does not proceed to the next period. With treatment 1, his expected utility is $.5 \times 1 + .25 \times 1 + .25 \times 1 = 1$, where $.25 = .5 \times .5$ is his probability of living through the second period. Treatment 2, similarly, offers a first-period probability of survival of .5. However, if the patient survives, the probability of surviving period 2 rises to .67. As a result, the expected payoff over 2 periods from treatment 2 is $.5 \times 1 + .5 \times .67 \times 1 + .5 \times .67 \times 1 = 1.17$, which is 17 percent greater than the expected rewards from treatment 1. These results are quite general both for situations of experimentation and adaptive behavior and for situations in which there is a sequence of lotteries on life and death (Berry and Viscusi 1981).

This preference-for-uncertainty result is apparent from a simple example. Suppose that the prior beliefs were based on a great deal of information, such as the equivalent of having observed 200 trials, of which 100 had seen success. In that instance, for treatment 2 the posterior probability of success in the second period after a successful outcome in the first period would be only $101/201 = .502$. The uncertain treatment consequently would offer a negligible advantage over treatment 1, whether the problem was one of costless learning and experimentation or a sequence of lotteries on life and death. If, however, there are many periods to the choice problem, the role of learning with the uncertain treatment 2 could potentially play an important role even with fairly tight prior risk beliefs. Trying the uncertain drug offers option value. In much the same

13. Two-armed bandit problems can become highly complex in multiperiod settings, particularly if both arms offer learning possibilities. In our example, if there were many periods, it would be worthwhile sticking with treatment 2 even if it looked inferior, given the benefit should it turn out to be superior.

way that asset volatility enhances the value of financial options, imprecise beliefs about risk enhance the potential gains from experimentation. In situations of experimentation and information acquisition, one would expect the uncertainty regarding the properties of a drug to be positively correlated with the degree to which it is possible to learn about the risks before the drug is on the market.

Our examples involved a single individual. If the FDA were to consider future users of the drug, the value of the learning that would come from trying the uncertain drug would multiply by many orders of magnitude. And note that this external benefit could be reaped even though, as ethics requires, early users would have to benefit in expectation from trying the drug.[14]

4.2. Learning from International Experiences

Drugs that are first adopted in other countries offer extensive experience from abroad in addition to the informational base of the US drug review procedures, which suggests that these drugs should be safer. However, contrary to expectations, the opposite result has proven true in practice. Drugs approved elsewhere before marketing in the United States have higher rates of adverse drug reactions than those first approved by the FDA (Olson 2013). This result is consistent with nonrandom sorting, in which drugs of uncertain riskiness get marketed elsewhere first.[15] Moreover, drugs posing few risks could receive accelerated approval by the FDA, which also would give a lower adverse reaction rate to drugs getting first approval in the United States. Other determinants of riskiness with respect to adverse drug reactions, hospitalizations, and deaths accord with expectations. For example, lengthier review periods decrease risks, and accelerated approvals increase risks.[16]

---

14. Thus, we are not proposing that A try the drug because the expected loss to him is less than the expected gain to B, C, and D.

15. Thus, it is the baseline rate of the risk of drugs approved elsewhere first, not FDA review practices, that produces this higher risk. This could be because drug companies disproportionately submit risky drugs elsewhere first. A second explanation would be that the FDA conducts a longer review than foreign countries when a drug looks risky, and the appearance of risk makes a drug riskier even if approved. A third explanation would be that the FDA simply counts the same level of foreign evidence of safety as being more compelling than its own evidence; that seems implausible.

16. These results would surely apply if there were no selection associated with the length of the review process. However, if drugs thought to be riskier got longer reviews, and those perceived as less risky got speedier approvals, we would expect a selection effect that pushed in the opposite direction from the observed pattern.

5. RISK AMBIGUITY AND BEHAVIORAL ANOMALIES THAT UNDERCUT RATIONAL CHOICE

A number of behavioral anomalies come into play when choices of drugs with uncertain properties are involved. These anomalies may be enshrined in government policies, both because government regulators and legislators are human and subject to behavioral biases and because the general public, which is subject to such behavioral phenomena, pressures the direction of government policy.

5.1. Ambiguity Aversion

Our analysis has assumed that the FDA knows the risks associated with a drug. However, probabilities are often highly uncertain. That is, there is ambiguity in the probability (for a broad discussion of ambiguity aversion, see Trautmann and van de Kuilen 2015).

Uncertain probabilities come from what is called a compound lottery. Thus, in effect, the FDA has to cast its vision forward and ask what the possible probabilities are that it will believe in the future, and what the likelihood of each is. To illustrate numerically, the FDA might think that the possible future probability assessments for a particular side effect are 10 percent, 5 percent, 1 percent, and .1 percent. Each of these would be associated with a set of future evidence. The FDA thinks that these probabilities will occur respectively with likelihoods .05, .05, .178, and .722. To get the expected probability, we compute the weighted average, for which each probability is weighted by its likelihood—for this example, .05 × 10 percent + .05 × 5 percent + .178 × 1 percent + .722 × .1 percent = 1.0002 percent.

Rational decision theory—see, for example, Raiffa (1968)—would prescribe that an individual facing a single-period choice between a drug with a certain 1 percent chance of side effect A should also be indifferent to taking the drug with an unknown risk of this side effect, but where the expected value was 1 percent, in a single-period decision. His chance of suffering from the side effect would be 1 percent in either case.

A particularly germane starting point for how people respond to ambiguity is the Ellsberg paradox (Ellsberg 1961). In an experimental task involving the chance to win a prize on the basis of draws from an urn, subjects prefer to make a choice from an urn presenting a known probability of success rather than from an urn involving an equivalent but uncertain probability of success, which generates a phenomenon known as ambiguity aversion.

There has long been speculation as to why we observe ambiguity aversion rather than neutrality with respect to uncertainty in single-period trials such as this. One possibility is that blame (including self-blame), recrimination, and regret may be greater when the subject fails to win the prize by making an unsuccessful choice from the uncertain urn after passing up a guaranteed 50 percent chance of success with the known probability (Heath and Tversky 1991). There also may be the fear that the experimenter could be manipulating the contents of the uncertain urn strategically to undermine the person's chance of success, but this risk can be neutralized by flipping a fair coin so that the uncertain option can always be transformed to being characterized by hard probabilities (Raiffa 1961).

The situation of ambiguity involving drugs is somewhat different since there is an ambiguous chance of losses rather than only gains as in the Ellsberg paradox study and most of its sequels. Evidence does exist that people often exhibit ambiguity aversion with respect to losses, but the level of ambiguity aversion appears to be less pronounced for losses than for gains (Camerer and Weber 1992).[17]

The power of ambiguity aversion for real-world decision makers is manifested by the responses of state judges to a product-liability scenario involving a drug with uncertain properties. The study asked 91 state court judges to choose between two different contrast agents then in use for computed tomography (CAT) scans (Viscusi 1999), although the probabilities presented to the judges were hypothetical. The judges considered the following text:

You are running a pharmaceutical company and must choose only one of two variants of a drug to market as a contrast agent for CAT scans. Your company has been selling *Old Drug* for decades. *Old Drug* works well, but there is a well-established 1/100,000 chance that the patient will suffer an adverse reaction and die. Your research team has already developed *New Drug* that performs identically and will have the same price and manufacturing costs, but *New Drug* poses uncertain risks. Based on the clinical trials the best estimate of the expected level of risk is that it is 1/150,000, but the risk is not known for sure. Some scientists believe the risk from *New Drug* is zero and some believe the risk could be as high as 1/50,000. Which drug would you choose to market? You must pick one. Circle your choice below.

17. Although the phenomena are different, this seems to complement the finding of prospect theory (Kahneman and Tversky 1979) that individuals are risk taking in the loss domain.

Although the new drug offered a lower expected risk level, 57 percent of the judges selected the old drug. After the survey, the judges were asked why they had selected the old drug. The dominant response was that a drug with known properties posed fewer liability risks and that they did not want to gamble with the uncertain drug as that would lead to greater potential liability if a patient died.[18] In effect, their comments suggest that ambiguity aversion on the part of jurors will lead the courts to make adverse rulings against uncertain drugs when products posing precisely known risks are compared with those whose risk levels are uncertain. This could be because errors of commission count much more than errors of omission. In addition, if the new drug is never tried, there is no way to know whether losses have been incurred, but with a switch to the new drug, any losses will be evident. Decision-maker preferences may stick with the well-chronicled preference for the status quo (Samuelson and Zeckhauser 1988). While most decision theorists would label this irrational, it is consistent with more general fears as to whether novel products incur greater expected liability costs than risky products posing more familiar, well-known hazards.

Consideration of the product-risk case illuminates another potential source of the aversion to ambiguous risks. For a product posing well-known risks, an adverse experience does not alter risk beliefs. For an uncertain product, however, an adverse event will increase the assessed risk. To the extent that people exhibit reference-point effects regarding risk levels, there will be an exaggerated response to adverse events that raise the assessed risk level.[19] The anticipation of such an exaggerated response in combination with the threat of liability will promote caution on the part of both companies seeking to market drugs and regulators who determine whether they can come on the market.

There is an important, almost philosophical, question as to whether the FDA should embrace this nonrational attitude. If it does not, it is effectively overruling consumer sovereignty. If it does, over the long run it will be accepting drug x that incurs a risk of k but turning down drug y that offers the same benefits but incurs a risk of $k - d$, where d is a positive number.

18. Note the similarity to the attitudes of the rational drug company in the example just considered.

19. There is, for example, a considerable asymmetry between consumers' willingness to pay for risk reductions and their willingness to accept amounts that they require for increases in risk (Viscusi, Magat, and Huber 1987).

For now, let us posit that the FDA followed the rational prescription and regulated on the basis of the expected value of risks. Additional considerations come into play. Consider the standpoint of consumers and then the standpoint of drug companies, both assumed to be rational for this analysis. Posit further that after the risks become known, the drug will be allowed to continue on the market if its side-effect risk is 1 percent or .1 percent but not if it is 10 percent or 5 percent. The a priori probabilities of the 10 percent and 5 percent risk are each .05. Thus, there is a 10 percent chance that the drug will be withdrawn from the market.

Until the risk becomes known, consumers experience an expected risk of 1 percent, so they are unaffected. But once the future information becomes known and the decision to continue or withdraw is made, they are better off than they would be with a certain 1 percent risk. To see this, consider what would happen if the FDA merely ignored the updated assessment. Consumers would then face a 1 percent expected risk forever, just the equal of the certain 1 percent risk. But the FDA conducts the analysis and determines that consumers are better off not taking this drug with 10 percent or 5 percent side-effect risks. Thus, consumer welfare is improved. Another way to see this is to observe that 90 percent of the time when the drug is left on the market, the consumer is better off, since more than four-fifths of the time the side-effect risk is only .1 percent and the other one-fifth of the time it is 1 percent as before.

What about the well-being of the companies? They would be better off with the certain 1 percent risk. First, their drug will stay on the market for sure. Second, they may be subject to liability if the side-effect risk is great enough that they are required to withdraw the drug from the market. Thus, there is not the danger that drug companies—positing that they are rational—would favor the uncertain drug. The regulators, from a purely parochial standpoint assuming that they are blame averse, would be better off with the certain risk. There is no chance of the embarrassment or criticism that comes from approving a drug and then removing or having the drug company remove it from the market because its risk is excessive.

Until a drug's performance is known for sure, there is always a positive informational externality going from current to future users. For example, consider a two-person world. Persons A and B both suffer from a condition for which drug 1 is known to perform at level 70, where the score balances efficacy and safety to produce a VN-M utility. Drug 2 has been tested, but its performance is ambiguous. The current analysis is

that its VN-M value is equally likely to be either 58 or 80 after one trial, which implies an expected utility of 69. After one person takes it, its future VN-M value will be known. Persons A and B have two strategies to compare. They could each take drug 1 and each receive a VN-M value of 70. Alternatively, they could flip a coin to decide who goes first. That individual would get a VN-M value of 68. The second person would stick with drug 2, if the first person had success, and otherwise would switch to drug 1. His VN-M would thus be $(.5)80 + (.5)70 = 75$. Following this second strategy, the expected VN-M value for A and B would be $(68 + 75)/2 = 71.5$, which represents an improvement.

This positive informational externality has been widely discussed, notably by those concerned with ethics. Individuals do not flip a coin to decide who goes first. The first patient is the first to take the drug. Having one person accept an expected loss so another, or thousands of others, can reap expected gains is contrary to contemporary medical ethics, which prescribe that the treatment be given that is best for this patient. Compensation, financial or other, could be employed to benefit the first taker. But this approach too, although welcomed by some, has generally been rejected.

At various points in history, drug trials have been undertaken with some form of compensation. For example, prisoners have been given time off in exchange for taking a drug that could be expected to hurt them. One of the most infamous and disgraceful clinical experiments in history was the Tuskegee syphilis study, whereby rural black men were left with untreated syphilis presumably because they were being given free medical care by the government. The compensation was clearly woefully inadequate. The Tuskegee patients were neither fully informed nor in any sense made better off by participating in the trial.

Our current drug regulatory system on occasion allows drugs with highly uncertain risks on the market, but that happens only when the condition the drug addresses is extremely debilitating or strongly life threatening. To be sure, if the compound lottery associated with an uncertain risk could be resolved in a day or a month, it would be worthwhile waiting. But often additional testing would be extremely expensive, and the only realistic way to know of the existence of risks, particularly low-probability risks, and the levels they impose is to allow the use of the drug and engage in postmarketing surveillance. It is sometimes said that if none of your bridges are falling down, you are building them too strong. Similarly, if none (or indeed very few) of the drugs the FDA has

approved are being withdrawn from the market, it is almost certainly regulating them too tightly.[20] Drug performance is highly uncertain. If our regulatory processes are avoiding all losses, then they are surely sacrificing significant gains.

We should note, however, that postmarketing surveillance is far from a perfectly rational process. Posit a drug already on the market that has proved more dangerous than expected and that would not have been allowed on the market in the first place given its currently assessed levels of efficacy and safety. In theory, it should be removed from the market. (We are leaving aside the consideration that patients should not quickly be taken off of a drug or might not replace a drug with another that is superior.) However, the FDA, like any agency, is hesitant to admit to past errors. Moreover, patients currently taking the drug may favor its continuation because status quo bias triumphs. Note the continued permission to use the diabetes drug Avandia in the United States despite considerable evidence and thousands of successful lawsuits deriving from its alleged association with a vast increase in heart attacks.[21]

Just over 3 decades ago, we conducted a decision-analysis training program for FDA officials. When the participants were asked why they did not want to base decisions on subjectively assessed risk levels when there was some uncertainty about the potential adverse effects of a drug, the dominant response was that they did not want to be responsible for the next thalidomide, a drug that led to severe birth defects among children born to mothers in Europe who had taken the drug for morning sickness. Interestingly, the morning-sickness drug thalidomide had never been approved by the FDA, but it remained one of the most salient cases in the regulatory literature. More interesting still, it has since been approved by the FDA and is currently marketed by Celgene as an anticancer treatment Thalomid.

---

20. See the Appendix for a discussion of how uncertainty in the testing regimen should combine with the FDA's prior distribution on the risk level to determine what cutoff in the drug trials should be used before market approval. Surprisingly, if the prior belief about safety is favorable (that is, the likelihood is great that a drug is safe), then the more uncertainty in the testing, the higher the cutoff—the less convincing the results—should be for allowing market approval.

21. An FDA panel refused to recommend that Avandia be withdrawn from the market, contrary to the recommendations of some of its reviewers. The *New York Times* reported its association with an increase of 83,000 heart attacks (Harris 2010). The use of Avandia has been substantially reduced since this information about its risks became known. The drug has been taken off the market in Europe.

5.2. Ambiguity-Seeking Behavior

Responses to uncertain probabilities may depend on the level of the probabilities. There could be manifestations of ambiguity aversion, ambiguity neutrality, and ambiguity-seeking behavior for the same type of risk, depending on the particular levels of risk involved. Ambiguity-seeking behavior is possible in lotteries involving low probabilities for gains and high probabilities for losses (Camerer and Weber 1992).

In a survey of 266 business owners and business managers in coastal North Carolina (Viscusi and Chesson 1999), the authors elicited preferences for locating a business when the choice involved either an area where the experts were in agreement about a precise risk or an area with experts in disagreement and thus holding different risk estimates. For low levels of average risk, subjects exhibited ambiguity aversion. At a risk of about .5, they exhibited ambiguity neutrality, and for higher risk levels they exhibited ambiguity-seeking behavior. Thus, risk ambiguity appeared to generate fear of the worst-case outcome at low probabilities and hope for the more favorable outcome at high probabilities.[22]

An analogous situation often arises in medical contexts. Consider two different scenarios for a doctor's visit. In scenario 1, a patient is told that there is a 10 percent chance that she has a fatal form of cancer, whereas in scenario 2, the patient is told that the chance could be as low as 5 percent or as high as 15 percent. The patient, reflecting ambiguity aversion for low-level risks, would prefer the known risk with scenario 1. If, however, scenario 1 poses a well-established 80 percent risk of a fatal form of cancer, whereas scenario 2 poses a risk that ranges from 70 percent to 90 percent, the same patient might find the medical uncertainty to be a hopeful sign of a better chance of survival.

5.3. Ambiguity and Risk Debates

How people react to ambiguous medical choices depends also on the source of the information that generates the ambiguity. Usually there is some disagreement either across studies or among different experts. For breast implants, the plaintiffs' bar and many medical experts, including the head of the FDA, either sounded the alarm or expressed caution, whereas plastic surgeons stood by the controversial technology. How do people react to situations in which there is conflicting information? In situations in which there are two conflicting studies of risk levels in

22. This seems consistent with risk taking on losses.

the ambiguity-averse range by a common entity, such as two government studies, people tend to average the risk assessments but with some modest additional premium that reflects ambiguity aversion. The ambiguity-aversion premium becomes more prominent if the studies are performed by different parties, such as a government entity and an industry group that has a vested interest in the research (Viscusi 1997). Whenever there is such a difference of opinion by different types of groups, people tend to gravitate to the upper-bound risk assessment value, irrespective of whether it is the industry or the government that argues for the high end of the range, which leads to a larger ambiguity-aversion premium.

Alarmist reactions of this type may contribute to the very strong market response to many risk controversies in the medical field, whether the controversies pertain to medical devices such as breast implants or to drugs such as Avandia. Such controversies are further fueled if they call people's attention to a novel risk not formerly associated with the product. The role of accustomed reference points in thinking about risks makes the perceived changes in risk levels associated with such debates a particular source of concern.

6. CONCLUSION

The FDA's regulatory approach to pharmaceuticals embodies many concerns that are clearly pertinent to a sound drug policy such as focusing on the reduced risk attributable to the drug, the increase in risk caused by the drug, and a balancing of these competing consequences. There is considerable evidence that pharmaceutical innovation has significantly improved health outcomes. However, despite this record of success, unexploited opportunities remain.

Even when the pertinent probabilities are known with precision, there is policy controversy as to how safety and efficacy should be appropriately balanced. Behavioral biases against errors of commission and asymmetric attitudes with respect to increases and decreases in risk will lead to inordinate attention to possible increases in risk.

The policy challenge becomes much greater when the risks of drugs are not well known. Despite the FDA's extensive testing requirements, the risks posed by a drug often cannot be well understood until the drug is put on the market, as evidenced by the frequency of recalls, labeling changes, and off-label uses of drugs for purposes for which the FDA has not assessed safety and efficacy. In addition, the presence of uncertainty

may prevent potentially effective drugs from reaching the market. The aversion to ambiguity, which is a fairly deep-seated phenomenon for low probabilities of loss, will disadvantage uncertain drugs as compared with drugs with similar, but more precisely understood, risks. But this bias against uncertainty is the opposite of what is optimal from the standpoint of promoting public health. Uncertainty about outcomes offers the only opportunity to learn about favorable properties of drugs and therefore to find new ways to promote long-term survival and improve the quality of life.

APPENDIX: DRUG SAFETY

A1. Model

Consider the public policy question of deciding how to set safety standards for drugs to be released on the market. We assume that there is one salient feature when deciding whether to release a drug: the probability that a patient experiences side effects when using the drug. Denote the threshold at which a drug will be released p_{thresh}; if the FDA believes that the drug has a probability of side effects $p \leq p_{\text{thresh}}$, it will release the drug to market and will reject the drug otherwise.

The FDA's belief about p (the unknown probability that the drug being considered will cause a side effect in a particular patient) depends on two things. First, on the basis of observations of previous drugs, the agency has a prior distribution P_{prior} over p; this is its prior belief about the likelihood of any randomly selected drug having a particular side-effect rate. Second, it will perform a clinical trial of the drug on a group of C people and observe the percentage of people who experience side effects. On the basis of this clinical trial, the FDA updates its prior belief to its posterior belief P_{post}, its belief about the likelihood that this particular drug will have a particular side-effect rate. Finally, if P_{post} has a mean greater than p_{thresh}, the FDA will reject the drug and will accept it otherwise.

A2. Beta Distribution

To model P_{prior}, we use the beta distribution. The beta distribution is a suitable choice because it is the conjugate prior belief of the binomial distribution. This family of distributions has several nice properties.

1. Given a binary experiment with an unknown p probability of success, a prior distribution over p of beta(α, β), and an observation of s successes and f failures, the posterior distribution on p is beta($\alpha + s$, $\beta + f$). This is what is meant when we say that the beta distribution is the conjugate prior of the binomial distribution.

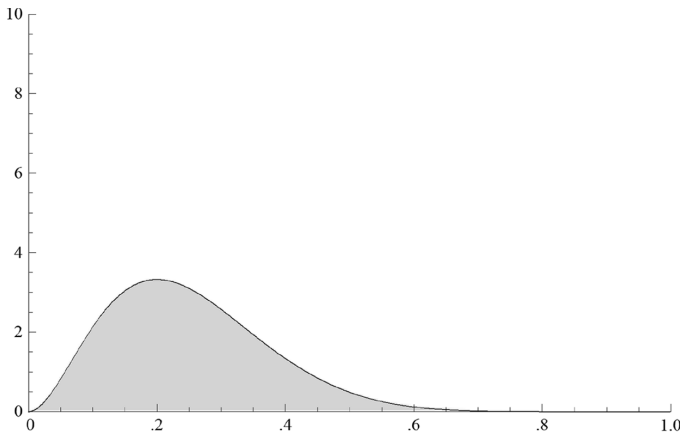2. Since beta(1, 1) is just the uniform distribution between [0, 1], we can inter-
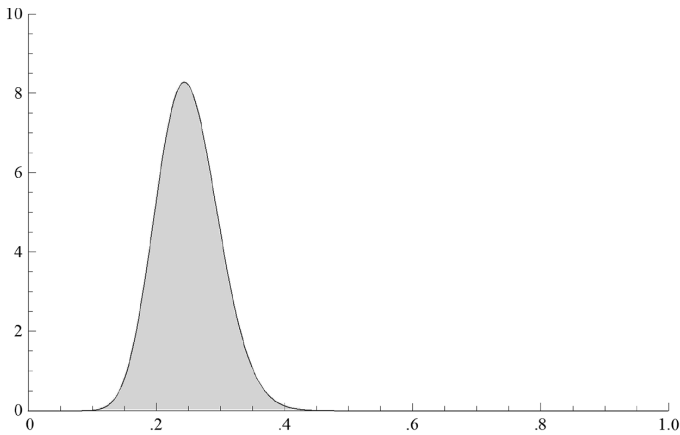
Figure A1. Prior belief of beta(3, 9)

Figure A2. Prior belief of beta(20, 60)

pret beta(α, β) as the correct posterior distribution after observing $\alpha - 1$ successes and $\beta - 1$ failures, given a uniform prior distribution.

3. Beta(α, β) has a mean of $\alpha/(\alpha + \beta)$.

4. Beta(α, β) has a mode of $(\alpha - 1)/[(\alpha - 1) + (\beta - 1)]$.

Finally, we note that given properties 2 and 4, for two beta distributions beta(α_1, β_1) and beta(α_2, β_2) where $\alpha_1/(\alpha_1 + \beta_1) = \alpha_2/(\alpha_2 + \beta_2)$ but $\alpha_1 + \beta_1 < \alpha_2 + \beta_2$, both beta distributions have the same mean, but the latter distribution expresses more confidence about that mean. For example, beta$(3, 9)$ is shown in Figure A1, and beta$(20, 60)$ is shown in Figure A2. In other words, since we can interpret the expression $\alpha + \beta - 2$ as the number of observations we have made of

our experiment, the higher this value, the more confident we are about the value of p, the rate of success. (A parallel analysis would apply for the mode.)

A3. Analysis

Given this model of approving drugs, we wish to determine the cutoff k that should be set for a clinical trial. Given a prior distribution $P_{\text{prior}} = \text{beta}(\alpha, \beta)$ and clinical trial of size C, we wish to determine the maximum percentage of patients in the trial who experience side effects (k) that will still lead to a posterior distribution P_{post} with a mean less than or equal to p_{thresh}. In other words, given our method for deciding whether to accept a drug as described above, we wish to determine what rates of side effects in our clinical trial are acceptable. We examine three different cases.

A3.1. P_{prior} Has a Mean Equal to p_{thresh}. In this case, $k = p_{\text{thresh}}$. To see this, we note that our prior $P_{\text{prior}} = \text{beta}(\alpha, \beta)$ has the property that $\alpha/(\alpha + \beta) = p_{\text{thresh}}$. Given a clinical trial with xC patients with side effects and $(1 - x)C$ patients with no side effects, our posterior $P_{\text{post}} = \text{beta}[\alpha + xC, \beta + (1 - x)C]$. Now consider the following:

$$\text{Mean}(P_{\text{post}}) = \frac{\alpha + xC}{\alpha + \beta + C} \le p_{\text{thresh}},$$

$$\alpha + xC \le p_{\text{thresh}}(\alpha + \beta + C),$$

$$\alpha + xC \le \alpha + p_{\text{thresh}}C,$$

$$x \le p_{\text{thresh}}.$$

In other words, the mean of our posterior is less than p_{thresh} if and only if $x \le p_{\text{thresh}}$. Accordingly we should set k equal to p_{thresh}.

A3.2. P_{prior} Has a Mean Less than P_{thresh}. In this case, $k > p_{\text{thresh}}$. To see this, we note that our prior $P_{\text{prior}} = \text{beta}(\alpha, \beta)$ has the property that $\alpha/(\alpha + \beta) < p_{\text{thresh}}$. For convenience, we set d equal to $p_{\text{thresh}} - \alpha/(\alpha + \beta)$.

Given a clinical trial with xC patients with side effects and $(1 - x)C$ patients with no side effects, our posterior $P_{\text{post}} = \text{beta}[\alpha + xC, \beta + (1 - x)C]$. Now consider the following:

$$\text{Mean}(P_{\text{post}}) = \frac{\alpha + xC}{\alpha + \beta + C} \le p_{\text{thresh}},$$

$$\alpha + xC \le p_{\text{thresh}}(\alpha + \beta + C),$$

$$\alpha + xC \le \alpha + d(\alpha + \beta) + p_{\text{thresh}}C,$$

$$x \le p_{\text{thresh}} + \frac{d(\alpha + \beta)}{C}.$$

In other words, the mean of our posterior is less than p_{thresh} if and only if $x \le p_{\text{thresh}} + [d(\alpha + \beta)]/C$. Accordingly, we should set k equal to $p_{\text{thresh}} + [d(\alpha + \beta)]/C$. Furthermore, since $d, C, \alpha, \beta > 0$, we have that $k > p_{\text{thresh}}$, as desired.

In addition, we can make two observations about k. The more skewed to the right of p_{thresh} our prior distribution is and the larger are d and $\alpha + \beta$, the larger is the difference between k and p_{thresh}. Furthermore, as $C \to \infty$, $k \to p_{thresh}$; the larger our clinical trial, the closer we should set k to p_{thresh}.

A3.3. P_{prior} Has a Mean Greater than P_{thresh}. In this case, $k < p_{thresh}$. To see this, we note that our prior $P_{prior} = \text{beta}(\alpha, \beta)$ has the property that $\alpha/(\alpha + \beta) > p_{thresh}$. For convenience, we set d equal to $\alpha/(\alpha + \beta) - p_{thresh}$.

Given a clinical trial with xC patients with side effects and $(1 - x)C$ patients with no side effects, our posterior $P_{post} = \text{beta}[\alpha + xC, \beta + (1 - x)C]$. Now consider the following:

$$\text{Mean}(P_{post}) = \frac{\alpha + xC}{\alpha + \beta + C} \leq p_{thresh},$$

$$\alpha + xC \leq p_{thresh}(\alpha + \beta + C),$$

$$\alpha + xC \leq \alpha - d(\alpha + \beta) + p_{thresh}C,$$

$$x \leq p_{thresh} - \frac{d(\alpha + \beta)}{C}.$$

In other words, the mean of our posterior is less than p_{thresh} if and only if $x \leq p_{thresh} - [d(\alpha + \beta)]/C$. Accordingly, we should set k equal to $p_{thresh} - [d(\alpha + \beta)]/C$. Furthermore, since $d, C, \alpha, \beta > 0$, we have that $k < p_{thresh}$, as desired.

In addition, we can make two observations about k. The more skewed to the left of p_{thresh} our prior distribution is and the larger are d and $\alpha + \beta$, the larger is the difference between k and p_{thresh}. Furthermore, as $C \to \infty$, $k \to p_{thresh}$; the larger our clinical trial, the closer we should set k to p_{thresh}.

A4. Modeling Error

To this point, we have assumed that our clinical trial is completely error free. However, in practice, there is some uncertainty to the observed results of a clinical trial; the true number of patients in the study who experience side effects may be higher or lower than the observed quantity. We model this situation as follows. Given a clinical trial of size C with xC observed side effects, we have some probability distribution over the true number of side effects. The only requirement we have about this distribution is that it is symmetric around xC. By symmetric, we mean the following. Let e_i denote the probability that the true number of side effects was $xC + i$. Then we have that

$$e_i = 0 \quad \text{for all } i < -xC, \ i \geq (1 - x)C,$$

where

$$e_i = e_{-i}.$$

It turns out that given a prior distribution of $P_{\text{prior}} = \text{beta}(\alpha, \beta)$, the mean of our posterior distribution P_{post} does not change after accounting for error; it is still $(\alpha + xC)/(\alpha + \beta + C)$. This means that our analysis about k—the maximum acceptable percentage of observed side effects in a clinical trial—from Section A3 does not change even after accounting for error. Regardless of the error distribution around the observed side-effect frequency, as long as it is symmetric, the same k is chosen given the same prior, clinical trial size, and p_{thresh}.

To see why the mean of the posterior distribution does not change, consider the following. We denote the probability density function (PDF) of a beta(x, y) distribution by $\beta_{x,y}$. Furthermore, we denote the PDF of our posterior distribution given a clinical trial of size C with xC observed side effects and a prior of beta(α, β) by f_{post}. Then we have

$$f_{\text{post}}(p) = \sum_{i=-xC}^{(1-x)C} e_i b_{(\alpha+xC+i),\,[\beta+(1-x)C-i]}$$

and

$$\text{Mean}(P_{\text{post}}) = \int_0^1 p f_{\text{post}}(p) dp$$

$$= \int_0^1 p \sum_{i=-xC}^{(1-x)C} e_i b_{(\alpha+xC+i),\,[\beta+(1-x)C-i]} dp$$

$$= \sum_{i=-xC}^{(1-x)C} e_i \int_0^1 p b_{(\alpha+xC+i),\,[\beta+(1-x)C-i]} dp$$

$$= \sum_{i=-xC}^{(1-x)C} e_i \text{Mean}\{\beta[\alpha + xC + i,\ \beta + (1-x)C - i]\}$$

$$= \sum_{i=-xC}^{(1-x)C} e_i \frac{\alpha + xC + i}{\alpha + \beta + C}.$$

However, since $e_i = e_{-i}$, we have that $\text{Mean}(P_{\text{post}}) = (\alpha + xC)/(\alpha + \beta + C)$, as desired.

A5. An Example

Here we work out a simple example to demonstrate some of the principles we derived above. Assume that we want to release all drugs that we believe have an average probability of side effects of less than 3 percent; we set p_{thresh} equal to 3. First, assume that we have a prior belief of beta(4, 196) (see Figure A3), which has a mean of 2 percent. If we run a clinical trial of size 100, our threshold for the number of side effects in the trial will be five patients; at most 5 percent of the clinical trial patients can experience side effects for the drug to be released. This demonstrates that an optimistic prior belief (one with a mean less than our threshold) results in a side-effect cutoff in our clinical trial that is higher than our threshold.

If we run a clinical trial of size 200, our threshold for the number of side ef-
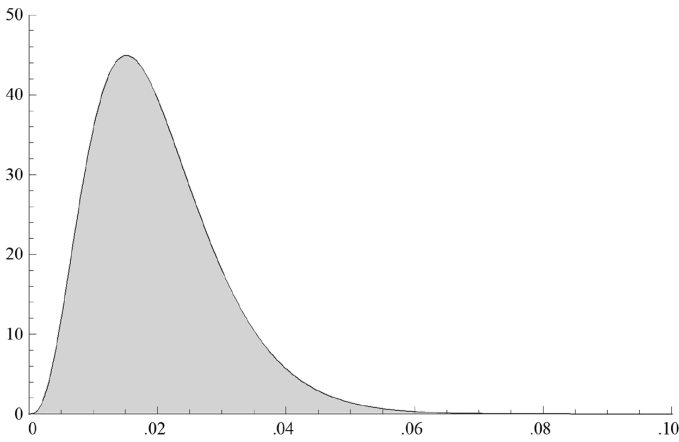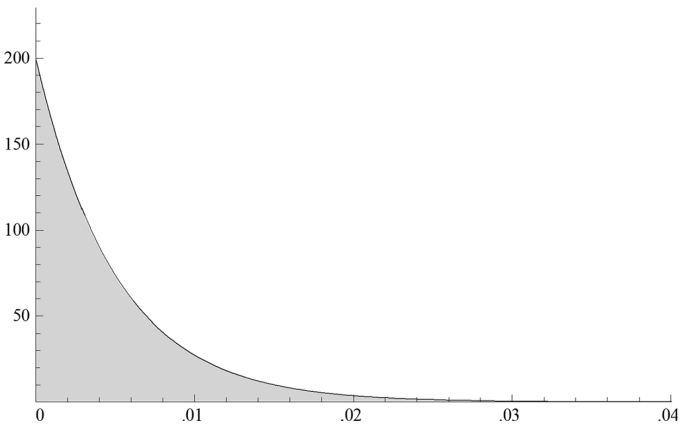
Figure A3. Prior belief of beta(4, 196)

Figure A4. Prior belief of beta(1, 199)

fects will instead be eight patients; at most 4 percent of the clinical trial patients can experience side effects for the drug to be released. This demonstrates that the larger the size of the clinical trial, the closer the cutoff percentage for side effects in the trial will be to the threshold.

Finally, if we instead had a prior belief of beta(1, 199) (see Figure A4), which has a mean of .5 percent and is significantly more skewed to the right, our thresholds for clinical trials of sizes 100 and 200 would be 8 percent and 5.5 percent, re-

spectively. This demonstrates that the more skewed our prior belief is to the right, the higher our cutoffs for clinical trials.

A6. Conclusion

Given this analysis, we come to the following conclusions about how clinical trial data should be interpreted. If an agency believes that drugs are generally safe and is very confident in that fact (P_{prior} is skewed to the left of p_{thresh}), then it should set its standards for clinical trial data to be higher than its actual desired probability of side effects. On the other hand, if the agency believes that drugs are generally dangerous and is very confident in that fact (P_{prior} is skewed to the right of p_{thresh}), then it should set its standards for clinical trial data to be lower than its actual desired probability of side effects. Finally, the larger the clinical trial (the greater C is), the closer the agency should set its standards for clinical trial data to its actual desired probability of side effects.

REFERENCES

American Society of Clinical Oncology. 2006. Reimbursement for Cancer Treatment: Coverage of Off-Label Drug Indications. *Journal of Clinical Oncology* 24:1–3.

Arnold, Matthew. 2010. AstraZeneca Settles Seroquel Off-Label Suit for $520m, CIA. *Medical Marketing and Media,* April 27. http://www.mmm-online.com /channel/astrazeneca-settles-seroquel-off-label-suit-for-520m-cia/article /168794/.

Begosh, Allan, John Goldsmith, Ed Hass, Randall W. Lutter, Clark Nardinelli, and John A. Vernon. 2006. Black Box Warnings and Drug Safety: Examining the Determinants and Timing of FDA Warning Labels. Working Paper No. 12803. National Bureau of Economic Research, Cambridge, MA.

Berry, Donald, and W. Kip Viscusi. 1981. Bernoulli Two-Armed Bandits with Geometric Termination. *Stochastic Processes and Their Applications* 11:35–45.

Brier, Jennifer. 2009. *Infectious Ideas: U.S. Political Responses to the AIDS Crisis.* Chapel Hill: University of North Carolina Press.

Camerer, Colin, and Martin Weber. 1992. Recent Developments in Modeling Preferences: Uncertainty and Ambiguity. *Journal of Risk and Uncertainty* 5:325–70.

Center for Medicare Advocacy. 2010. CMA Report: Medicare Coverage for Off-Label Drug Use. September 16. http://www.medicareadvocacy.org/cma-report -medicare-coverage-for-off-label-drug-use/.

Cook, Daniel M., Rama K. Gurugubelli, and Lisa A. Bero. 2009. Risk Management Policy and Black-Box Warnings: A Qualitative Analysis of U.S. FDA Pro-

ceedings. *Drug Safety* 32:1057–66.

Dickey, Robert W. 2012. FDA Risk Assessment of Seafood Contamination after the BP Oil Spill. *Environmental Health Perspectives* 120:a54–a55.

Downing, Nicholas S., Jenerius A. Aminawung, Nilay D. Shah, Joel B. Braunstein, Harlan M. Krumholz, and Joseph S. Ross. 2012. Regulatory Review of Novel Therapeutics: Comparison of Three Regulatory Agencies. *New England Journal of Medicine* 366:2284–93.

Edwards, Hayes. 2012. Risky Business: How the FDA Overstepped Its Bounds by Limiting Patient Access to Experimental Drugs. *George Mason University Civil Rights Law Journal* 22:389–426.

Ellsberg, Daniel. 1961. Risk, Ambiguity, and the Savage Axioms. *Quarterly Journal of Economics* 75:643–69.

FDA (Food and Drug Administration). 2002. *FY 2001 Performance Report to Congress for the Prescription Drug User Fee Act of 1992 as Reauthorized and Amended by the Food and Drug Administration Modernization Act of 1997*. Washington, DC: FDA. http://www.fda.gov/downloads/aboutfda/reportsmanuals forms/reports/userfeereports/performancereports/UCM135718.pdf.

———. 2007. *FY 2006 Performance Report to the President and Congress for the Prescription Drug User Fee Act*. Washington, DC: FDA. http://www.fda .gov/downloads/aboutfda/reportsmanualsforms/reports/userfeereports /performancereports/ucm103550.pdf.

———. 2011. Guidance for Industry Responding to Unsolicited Requests for Off-Label Information about Prescription Drugs and Medical Devices. http:// www.fda.gov/downloads/drugs/guidancecomplianceregulatoryinformation /guidances/ucm285145.pdf.

Fiegl, Charles. 2011. Medicare Ordered to Pay for Off-Label Drugs. *American Medical News,* March 18. http://www.ama-assn.org/amednews/2011/03/14 /gvse0318.htm.

Grady, Denise. 2013. Panel's Report Urges Breast Cancer Drugs for Healthy Higher-Risk Women. *New York Times,* April 16.

Hamilton, James T., and W. Kip Viscusi. 1999. *Calculating Risks? The Spatial and Political Dimensions of Hazardous Waste Policy*. Regulation of Economic Activity Series No. 21. Cambridge, MA: MIT Press.

Harris, Gardiner. 2010. Research Ties Diabetes Drug to Heart Woes. *New York Times,* February 19. http://www.nytimes.com/2010/02/20/health/policy/20 avandia.html.

Heath, Chip, and Amos Tversky. 1991. Preference and Belief: Ambiguity and Competence in Choice under Uncertainty. *Journal of Risk and Uncertainty* 4:5–28.

Kahneman, Daniel, and Amos Tversky. 1979. Prospect Theory: An Analysis of Decision under Risk. *Econometrica* 47:263–91.

Kolata, Gina. 2013. F.D.A. Plans to Loosen Rules on Alzheimer's Drug Approval.

New York Times, March 14.

Kozauer, Nicholas, and Russell Katz. 2013. Regulatory Innovation and Drug Development for Early-Stage Alzheimer's Disease. *New England Journal of Medicine* 368:1169–71.

Lammi, Glenn G. 2012. State Attorneys General Step to the Fore on Off-Label Drug "Promotion." *Forbes.com,* September 7. http://www.forbes.com/sites/wlf/2012/09/07/state-attorneys-general-step-to-the-fore-on-off-label-drug-promotion.

Lasser, Karen E., Paul D. Allen, Steffie J. Woolhandler, David U. Himmelstein, Sidney M. Wolfe, and David H. Bor. 2002. Timing of New Black Box Warnings and Withdrawals for Prescription Medications. *Journal of American Medical Association* 287:2215–20.

Lichtenberg, Frank R. 2011. The Quality of Medical Care, Behavioral Risk Factors, and Longevity Growth. *International Journal of Health Care Finance and Economics* 11:1–34.

———. 2012. The Effect of Pharmaceutical Innovation on the Functional Limitations of Elderly Americans: Evidence from the 2004 National Nursing Home Survey. *Advances in Health Economics and Health Services Research* 23:73–101.

———. 2013. The Effect of Pharmaceutical Innovation on Longevity: Patient Level Evidence from the 1996–2002 Medical Expenditure Panel Survey and Linked Mortality Public-Use Files. *Forum for Health Economics and Policy* 16:1–33.

———. 2014a. Has Medical Innovation Reduced Cancer Mortality? *CESifo Economic Studies* 60:133–77.

———. 2014b. Pharmaceutical Innovation and Longevity Growth in 30 Developing and High-Income Countries, 2000–2009. *Health Policy and Technology* 3:36–58.

Lipworth, Wendy, Ian Kerridge, Miles Little, and Richard Day. 2012. Evidence and Desperation in Off-Label Prescribing: Recombinant Factor VIIa. *British Medical Journal* 344:26–28.

Logan, Aaron C., Veronica Yank, and Randall S. Stafford. 2011. Off-Label Use of Recombinant Factor VIIa in U.S. Hospitals: Analysis of Hospital Records. *Annals of Internal Medicine* 154:516–22.

Miller, Kelli. 2009. Off-Label Drug Use: What You Need to Know. *WebMD.* http://www.webmd.com/a-to-z-guides/features/off-label-drug-use-what-you-need-to-know.

Moore, Thomas J., and Curt Furberg. 2012. The Safety Risks of Innovation: The FDA's Expedited Drug Development Pathway. *Journal of the American Medical Association* 380:869–70.

Murphy, Shirley, and Rosemary Roberts. 2006. "Black Box" 101: How the Food and Drug Administration Evaluates, Communicates, and Manages Drug Bene-

fit/Risk. *Journal of Allergy and Clinical Immunology* 117:34–39.

New York Times. 2013. Drugs for Early-Stage Alzheimer's. March 18. http://www.nytimes.com/2013/03/18/opinion/drugs-for-early-stage-alzheimers.html.

Nichols, Albert L., and Richard J. Zeckhauser. 1986. The Perils of Prudence: How Conservative Risk Assessments Distort Regulation. *Regulation* 10:13–24.

Olson, Mary K. 2013. Eliminating the U.S. Drug Lag: Implications for Drug Safety. *Journal of Risk and Uncertainty* 47:1–30.

Pollack, Andrew, and Mosi Secret. 2012. Amgen Agrees to Pay $762 Million for Marketing Anemia Drug for Off-Label Use. *New York Times,* December 18. http://www.nytimes.com/2012/12/19/business/amgen-agrees-to-pay-762-million-in-drug-case.html.

Radley, David C., Stan N. Finkelstein, and Randall S. Stafford. 2006. Off-Label Prescribing among Office-Based Physicians. *Archives of Internal Medicine* 166:1021–26.

Raiffa, Howard. 1961. Risk, Ambiguity, and the Savage Axioms: Comment. *Quarterly Journal of Economics* 75:690–94.

———. 1968. *Decision Analysis*. Reading, MA: Addison-Wesley.

Samuelson, William, and Richard Zeckhauser. 1988. Status Quo Bias in Decision Making. *Journal of Risk and Uncertainty* 1:7–59.

Schmit, Julie. 2004. Drugmaker Admitted Fraud, but Sales Flourish. *USA Today*, August 16.

Stohr, Greg. 2008. U.S. Supreme Court Refuses to Expand Experimental-Drug Access. *Bloomberg.com,* January 14. http://www.bloomberg.com/apps/news?pid=newsarchive&sid=atnVIGEgRQt4.

Thaul, Susan. 2012. How FDA Approves Drugs and Regulates Their Safety and Effectiveness. CRS Report No. 7-5700. June 25. Congressional Research Service, Washington, DC.

Trautmann, Stefan T., and Gijs van de Kuilen. 2015. Ambiguity Attitudes. Pp. 1:89–116 in *The Wiley Blackwell Handbook of Judgment and Decision Making,* edited by Gideon Keren and George Wu. Chichester: John Wiley & Sons.

Trautmann, Stefan T., and Richard J. Zeckhauser. 2013. Shunning Uncertainty: The Neglect of Learning Opportunities. *Games and Economic Behavior* 79:44–55.

Trowbridge, Ronald L., and Steven Walker. 2007. The FDA's Deadly Track Record. *Wall Street Journal,* August 14. http://www.wsj.com/articles/SB118705547735996773.

US General Accounting Office. 1980. *FDA Drug Approval: A Lengthy Process That Delays the Availability of Important New Drugs*. Report No. HRD-80-64. May 28. Washington, DC: US General Accounting Office.

Viscusi, W. Kip. 1997. Alarmist Decisions with Divergent Risk Information. *Economic Journal* 107:1657–70.

———. 1999. How Do Judges Think about Risk? *American Law and Economics*

Review 1:26–62.

Viscusi, W. Kip, and Harrell Chesson. 1999. Hopes and Fears: The Conflicting Effects of Risk Ambiguity. *Theory and Decision* 47:153–78.

Viscusi, W. Kip, Wesley A. Magat, and Joel Huber. 1987. An Investigation of the Rationality of Consumer Valuations of Multiple Health Risks. *RAND Journal of Economics* 18:465–79.

Viscusi, W. Kip, Wesley A. Magat, and Robert Scharff. 1996. Asymmetric Assessments in Valuing Pharmaceutical Risks. *Medical Care* 34:DS34–DS47.

Wang, Bo, Joshua J. Gagne, and Niteesh K. Choudry. 2012. The Epidemiology of Drug Recalls in the United States. *Archives of Internal Medicine* 172:1109–10.

Wilson, Duff. 2010. Novartis Settles Off-Label Marketing Case Over 6 Drugs for $422.5 Million. *New York Times*, September 30. http://www.nytimes.com/2010/10/01/health/policy/01novartis.html.

Yank, Veronica, C. Vaughan Tuohy, Aaron C. Logan, Dena M. Bravata, Kristan Staudenmayer, Robin Eisenhut, Vandana Sundaram, Donal McMahon, Ingram Olkin, Kathryn M. McDonald, Douglas K. Owens, and Randall S. Stafford. 2011. Systematic Review: Benefits and Harms of In-Hospital Use of Recombinant Factor VIIa for Off-Label Indications. *Annals of Internal Medicine* 154:529–40.

Zeckhauser, Richard J. 2010. Investing in the Unknown and Unknowable. Pp. 304–46 in *The Known, the Unknown, and the Unknowable in Financial Risk Management: Measurement and Theory Advancing Practice,* edited by Francis X. Diebold, Neil A. Doherty, and Richard J. Herring. Princeton, NJ: Princeton University Press.