

**“BRIGHT LINES” AND THE VALUE OF LIFE:
RESOLVING THE DISPUTE OVER
THE REGULATION OF CARCINOGENS**

JANICE C. WRIGHT

E-97-04

JUNE 1997

CITATION

This document appears as Discussion Paper 97-04 of Belfer Center for Science and International Affairs and as contribution E-97-04 of the Center's Environment and Natural Resources Program. Comments are welcome and may be directed to the author at the Belfer Center.

This paper may be cited as Janice C. Wright. "'Bright Lines' and the Value of Life: Resolving the Dispute over the Regulation of Carcinogens." BCSIA Discussion Paper 97-04, ENRP Discussion Paper E-97-04, Kennedy School of Government, Harvard University, November 1997.

The views expressed in this paper are those of the author and publication does not imply their endorsement by ENRP, BCSIA or Harvard University. Reproduction of this paper may only be made with the permission of the Belfer Center for Science and International Affairs, Publications, 79 JFK Street, Cambridge, MA 02138.

Foreword

Janice C. Wright was recipient of the 1996-1997 Joseph Crump Fellowship and received her Ph.D. from Harvard University in June 1997. This paper is a chapter of her dissertation, *Investments that Save Lives: The Norms of Environmental and Medical Decision Making*.

The Joseph Crump Fellowship, administered by the John F. Kennedy School's Center for Business and Government and Belfer Center for Science and International Affairs, is awarded annually to a doctoral candidate conducting research on the environment, natural resource management, or the intersection between energy and the environment.

**“Bright Lines” and the Value of Life:
Resolving the Dispute over the Regulation of Carcinogens**

Janice C. Wright

Introduction

The process of carcinogen regulation in the United States is characterized by conflict. One of these conflicts, that between the Office of Management & Budget (OMB) and the agencies that regulate carcinogen emissions, is stylized in this paper as a dispute over the choice of decision heuristic that should be used in the design of carcinogen standards. This dispute is framed as a problem in negotiation analysis, to see whether this may help resolve the conflict. It is shown that the formalization of the heuristic used by the regulatory agencies -- a bright line of acceptable individual risk -- leads to outcomes that appear to be at odds with the concerns of these agencies. Moreover, this heuristic is an accident of history, and lacks a theoretical rationale. A search for an improved heuristic that will capture the concerns of the regulatory agencies leads back to a decision-theoretic value-of-life model, developed nearly three decades ago. It is argued that this model captures the concerns of both parties in this dispute, the OMB's concern with allocative efficiency, and the regulatory agencies' concern with the distribution of risk across individuals. A simpler variant of this “value-of-life” model is developed, and leads to a surprising result.

Regulation of carcinogenic emissions and residues from production processes is a priority of certain regulatory agencies in the United States. The Food & Drug Administration (FDA), the Occupational Safety & Health Administration (OSHA), and the Environmental Protection Agency (EPA) are the three major federal agencies charged with reducing exposure of consumers, workers, and residents to chemical carcinogens.

The rational economist would prescribe a model of carcinogen regulation in which only rules with benefits exceeding costs would be promulgated. The application of this model requires the monetary valuation of statistical life, or more precisely, the valuation of an averted statistical case of cancer. While exact valuations of statistical life will remain elusive, welfare economists argue that five to ten million dollars is about right. Yet the cost of saving a life exceeds ten million dollars for many rules, sometimes by orders of magnitude, a situation of concern to economists at the Office of Management & Budget. However, most of the statutes governing carcinogen regulation require that regulatory agencies consider only the benefits of such regulations, and the courts have supported this interpretation. Further, the absence of "regulatory budgets" for private sector compliance costs means that agencies have little incentive for adopting economic efficiency as a decision heuristic, because the costs of compliance are borne by producers and consumers.

The decision heuristic that has been adopted by the regulatory agencies is that "significant" risks to individuals from each carcinogen are to be identified, and, where technically feasible, reduced to an acceptable level. Benchmarks of acceptable risk are known as *bright lines*. For instance, OSHA has set its bright line at an additional lifetime risk to the most exposed worker of a one in a thousand chance of developing cancer.

Two types of bright lines of individual risk have been defined in the courts. Risks above the *de manifestis* level are considered so large that they must be reduced; risks below the *de minimis* level are considered negligible, and are to be ignored. OSHA's bright line of one in a thousand is of the first type.

The notion of acceptable risk is ethically appealing, since the distribution of a risk burden across an exposed group matters, as well as its average magnitude. The elimination a mortality risk of 0.01 borne by 100 people, and the elimination of a mortality risk of 0.00001 borne by 100,000 people, will, on average, both "save" one life, but most of us would value the former opportunity more highly than the latter. Economists recognize this, but argue that the additional risk to individuals is so small in most cases of carcinogen exposure that the distribution issue can be safely ignored. However, how small is small enough remains an open question. A major shortcoming of the bright-line heuristic is its failure to distinguish between degrees of unacceptability of individual risk; it acts as a switch, and all increments in risk above the *de manifestis* bright line are regarded as equally unacceptable.

This paper begins with an explicit acknowledgement of the conflict that exists between the OMB and the regulatory agencies. Conflicts can sometimes be resolved through skillful negotiation by a third party. In this paper, the "*art and science of negotiation*", as developed by Howard Raiffa, is used to provide a framework for improving carcinogen regulation.¹ There are three steps needed for this type of negotiated settlement.

¹ Raiffa, 1982; 1996.

The *first* step is the identification of many possible outcomes in a joint “brainstorming” session, and the creative design of potential contracts between the disputing parties. The *second* step is the identification of the subset of the contracts that are efficient in the Pareto sense; the Pareto frontier, along which these contracts lie, is known as the contract curve. In the majority of real negotiations, argues Raiffa, the chosen contract lies far within this frontier, leaving gains for all parties on the table. Therefore, a major improvement would be possible through tracing out the efficient frontier, where win-win changes can occur. The *third* step is the negotiation of the choice of contract on the frontier.

Using the Raiffa framework of negotiation analysis, I turn the existing model of carcinogen regulation back to one of allocation of a “regulatory budget”. An annual regulatory budget would set a limit on the costs that an agency could impose on the private sector each year.² It is a two-party model. Party A is concerned with allocative efficiency and wishes to save as many lives as possible. This party could be thought of as the OMB. Party B is concerned with the distribution of risk and wishes to move as many exposed individuals as possible below the *de manifestis* bright line. This party could be thought of as FDA or OSHA or EPA.

Currently, agencies regulate carcinogens one-at-a-time, and uniform standards, that is, the same permissible maximum concentration, or emission rate, of the carcinogen for all exposed groups, are the norm. The first step of the negotiation model identifies possible outcomes as differentiated standards for a set of several carcinogens. In a differentiated standard, the permissible concentration of each carcinogen is allowed to vary across the

² A joint constraint on both parties is required for this type of negotiation analysis.

exposed groups. A set of hypothetical data is used to illustrate the second step, the generation of the contract curve. The uniform standard lies inside this contract curve, demonstrating that it does not maximize the objective of either party. The third step -- the negotiation along the frontier -- turns out to be unnecessary.

The generation of the contract curve requires the formalization of the objective functions of the two parties, and it is shown that the bright-line heuristic of the regulatory agencies logically formalizes into an objective function leading to outcomes that are manifestly unfair. Moreover, the bright-line heuristic has evolved in response to court interpretations of legislation, and, consequently, it lacks an underlying rationale. Such rationales that have been proposed are after-the-fact justifications.

Since I interpret the bright-line heuristic as reflecting a concern of the regulatory agencies for the distribution of risk over individuals, I next search for a heuristic that will capture this concern for fairness, but that is grounded in a theory. A prescriptive model of carcinogen regulation, in which the objective would be to maximize the number of saved lives, adjusted for fairness in the distribution of risk across individuals, would meet this requirement. Because such an objective would be both efficient (in the allocative sense), and fair (by weighting for the distribution of the burden of the risk), it would capture the concerns of both parties.

If a decision heuristic can be constructed that should satisfy both parties, then there is no longer a need for the third step of the Raiffa approach, the negotiation along the frontier. This is because the conflict over the allocation of a regulatory budget would be resolved, at least, in theory. Other conflicts between the parties would not be resolved; one could

envisage the dispute shifting to the design and levels of regulatory budgets.

A candidate for a new (efficient and fair) decision heuristic is found in a decision-analytic model that is known as the Black Pill problem.³ In the Black Pill problem, an individual is faced with a choice between accepting or rejecting a risk of death (the Black Pill) in return for monetary compensation. The amount of compensation rises more steeply than the risk; it follows that the value implicitly placed by the individual on his or her life also rises. This leads to the derivation of the value of life as a function of the magnitude of the individual risk, forming a basis for weighting the value of an individual's life in accordance with the size of the risk burden.

Although the Black Pill problem is analyzed from the perspective of an individual, the solution can be extended to the perspective of a group, whose members are exposed to risks of varying sizes. The value of life for a group with a given risk distribution can be determined from the risk-adjusted values of life for the individuals in the group. The underlying rationale of fair compensation for the individual is extended to the social setting by use of the Kaldor-Hicks criterion; that is, it should be possible to pay fair compensation to the risk-bearers.

A simple variant of the Black Pill problem is constructed, and a surprisingly extreme result emerges. The question then arises of whether this result is a consequence of the use of the exponential utility function (denoting constant risk aversion) in the model. Utility functions that represent a more realistic attitude to risk, decreasing risk aversion, are examined. The result holds, and a radical carcinogen regulation policy is presented.

³ Howard, 1980.

The Development of a Conflict

The statutory origin of carcinogen regulation in the United States is to be found in the Delaney Clauses, similarly worded clauses in three laws governing the FDA's use of risk assessment. The Delaney Clauses state that nothing that induces cancer in man or animal can be deemed safe. In the early 1970's, it became clear that the increase in ability to detect man-made carcinogens at very low levels made literal interpretation of the Delaney Clauses impossible. FDA then chose an additional lifetime risk of one in a million as "*essentially zero*", giving rise to one type of bright line -- a *de minimis* or negligible increment in individual lifetime risk.⁴

EPA has bright lines of two types for categorizing increments in individual risk. The *de minimis* level follows FDA -- as many people as possible should be below the one in a million level. However, because such protection is often technically infeasible, EPA has set a *de manifestis* bright line at one in ten thousand for the most exposed individual. As mentioned in the introduction, OSHA has set its *de manifestis* bright line at one in a thousand. A risk above a *de minimis* bright line is undesirable; a risk above a *de manifestis* bright line is unacceptable. Note that these bright lines refer to the additional *lifetime* risk to the *most exposed individual* from a *single carcinogen*, and a *single source*.

In 1981, OMB was given presidential authority to review federal regulations. Several case studies of carcinogen regulations commissioned at that time showed great variation in

⁴ Graham, 1995, p.21-22. The three laws containing the Delaney Clauses were passed in 1958, 1960, and 1968.

the cost of saving a life across industries and, within industries, across plants.⁵ When the stringency of an emission standard increases, the marginal cost of saving a life rises steeply. In one study of EPA's proposed standards for arsenic and radionuclides, the author estimates that 96% of the benefit that would come from the uniform standards proposed for these two toxins, could be gained at only 27% of the cost.⁶ These studies all point to the efficiency gains from setting *differentiated standards*, that is, varying the stringency of standards across industries or plants.

Following these studies, an ongoing interagency struggle ensued between the regulatory agencies and the OMB; from time to time the latter publishes a table showing new rules ranked by increasing cost-per-life ratios.⁷ The OMB has prevailed to some extent, with every president, from Reagan onward, ordering that costs and benefits for significant rules be identified in Regulatory Impact Analyses (RIAs).⁸ Although economists within the regulatory agencies take two to ten million dollars as the appropriate value of an averted case

⁵ Broder & Morrall, 1983; DeMuth, 1983; Haigh *et al.*, 1984.

⁶ DeMuth, 1983. DeMuth concludes by proposing two bright lines as a threshold for initiating regulatory action - a population risk of one death per year, and an individual risk of one in ten thousand - presumably believing that this would prevent the saving of very expensive statistical lives.

⁷ Morrall, 1986, contains the first of these tables. The OMB changes the calculations of the regulatory agencies to standardize the methods, but does not reveal how this is done (R. Belzer, pers. comm., Oct., 1995). There are two changes to regulatory agency practice that would make an enormous difference to estimations of the cost of saving a life. The first is the calculation of incremental, rather than average cost-effectiveness. The second is the discounting of the lives saved to allow for the latency period of carcinogenesis.

⁸ Currently, a "significant" rule is one in which compliance costs exceed more than \$100 million. Note that *identifying* costs and benefits is not the same as *balancing* costs and benefits.

of cancer⁹, the agencies have been forbidden by the courts to use cost-benefit analysis for setting standards, unless it is required by the relevant statute.¹⁰

In the 1990's, political pressure has grown for the use of benefit-cost analysis and for the setting of regulatory budgets to limit the costs that agencies can impose on the private sector each year. A regulatory budget is *"attractive in concept, but we are a long way from being able to put such a measure into practice."*¹¹ Nonetheless, application of the concept of a regulatory budget might be illuminating.

*"... the budget analogy is useful in highlighting the problems with the current regulatory effort. It should be equally helpful in designing solutions to those problems. In that sense, the concept of a regulatory budget may be viewed as an ideal type, within which a set of reforms can be molded. Understanding the ideal design may help sort out more limited, but also more practical, plans that may serve as a transition between the immediate and the ideal."*¹²

In summary, the reviewer of the rule-makers (OMB) wants differentiated standards and cost-benefit analysis, pointing to the enormous variation in the cost of averting a case of cancer. The rule-makers (FDA, EPA, OSHA) want to retain uniform standards and equal protection across individuals.

⁹ Fisher, *et al.*, 1989.

¹⁰ One key court challenge was a case involving OSHA's cotton-dust standard in 1981; another involved EPA's proposed standard for vinyl chloride in 1987. In both instances, the agencies were instructed to ignore the costs of compliance. Graham, 1995, pp.26-29; Van Houtven & Cropper, 1996.

¹¹ Portney, 1995, p.23.

¹² Litan & Nordhaus, 1983, p.133.

A Negotiation Model May Help Resolve the Conflict

There are then two different entrenched views on the choice of decision heuristic for the regulation of carcinogens. Since this is the case, and since regulations are always the result of negotiations, it may be that some resolution of the debate is possible through framing it as a two-party negotiation problem. We can denote the OMB as party A: an advocate of allocative efficiency. We can denote the regulatory agencies, such as FDA, EPA, and OSHA, as party B: an advocate of the protection of individuals.

Dispute resolution is a rapidly growing field, characterized by Raiffa as a blend of art and science.¹³ The art of negotiation “... *includes interpersonal skills, the ability to convince and be convinced, the ability to employ a basketfull of bargaining ploys, and the wisdom to know when and how to use them.*”¹⁴ The analysis presented in this paper is not concerned with the art of negotiation; rather it is focused on one piece of the science of negotiation -- the systematic analysis required to generate a “contract curve”, as developed by Raiffa.¹⁵ This approach has been developed within the Pareto framework of neoclassical economics, and a contract curve is a special type of production possibility frontier.

As summarized in the introduction to this paper, a simple application of Raiffa's

¹³ Fisher, *et al.*, 1991, is a primer for practitioners that blends the art and science of negotiation.

¹⁴ Raiffa, 1982, p.8.

¹⁵ The approach adopted in this paper is based on three lectures on negotiation analysis given by Professor Howard Raiffa at the Kennedy School of Government, Harvard University in Spring, 1996.

approach comprises three major steps.

The first step involves the creative design of possible contracts; a “template” listing the issues under dispute, alongside the options for resolution, is prepared. It is assumed in this paper, that what Raiffa terms the FOTE condition holds, that is, full, open, and truthful exchange of information; there is no strategic withholding of information, and both parties have unrestricted access to the same data. In the template for the carcinogen regulation problem, an issue is the exposure of a particular group to a particular carcinogen, and the options for resolution are a set of exposure limits, including the *status quo*. In other words, it is assumed that it is possible to set differentiated standards -- the exposure limit for a given carcinogen is allowed to vary across industries. Any number of carcinogens, and any number of emission standards, can be considered in one application of the model.

The second step is the generation of the contract curve. In Raiffa’s approach, the two parties begin by simply ranking the possible contracts. These ordinal rankings can be converted to cardinal rankings, by introducing a scoring system to allow both parties to express the strength of their preferences. Each party’s ranking of each potential contract is plotted on an axis. The usual graphical representation of Pareto efficiency follows; a contract is dominated if another lies to the northeast.

The ability to generate a sequence of Pareto-efficient contracts -- a contract curve -- is a major achievement of the “science” of negotiation. Empirically, the great majority of solutions to real conflicts lie far inside the contract curve.¹⁶ Processes for conflict resolution that are believed to be fair (such as coin tossing) generate solutions that are inefficient, and

¹⁶ Raiffa, 1996.

so gains are left behind on the negotiating table.¹⁷

In this paper, a contract curve is generated for a set of hypothetical proposed carcinogen regulations. Both parties begin by ordering their stylized preferences. Party A wishes to avoid over-valuation of statistical lives, and is indifferent to the distribution of the risk burden. Party B adopts a logical formalization of the bright line heuristic. No contracts -- agreements on sets of differentiated standards -- can be designed until both parties are subject to a joint constraint. Such a constraint is introduced in the form of a "regulatory budget" -- a limit on the amount that can be spent in reducing exposure to the carcinogens.

The third step in Raiffa's approach is the choice of contract from the set on the frontier. The contracts that lie at the extremes of the frontier are manifestly unfair, since they each represent the preferences of only one party. A process of negotiation should lead to one of the contracts away from the extremes. Particular concepts of equity can lead to game-theoretic formal solution of negotiation problems, but this is outside the scope of the analysis in this essay. Note though, that since each party experiences decreasing returns as it moves away from its starting position, a contract curve will be concave (bow outward), and the set of feasible contracts will be convex. Because of this convexity, a solution in the central

¹⁷ A procedurally fair solution like coin tossing does not allow for variation in the strength of individuals' preferences. Raiffa (1996) illustrates this point with the story of two students who each had a single free ticket to a basketball game, but each wanted two tickets. They tossed a coin to solve the problem, leaving one student with two tickets and the other with none. An alternative would be for both students to tell a facilitator how much he or she would be willing to pay for the second ticket. The student with the stronger preference (willing to pay a higher price) would purchase the ticket from the other at the average of the two selling prices, thus splitting the difference. Now one student would be left with two tickets and the other with some cash. This alternative procedure would lead to an outcome that is a Pareto improvement, and more equitable.

portion will give each party, say, seventy percent of their optimum, rather than fifty percent; in other words, the game is not zero-sum.

Figure 1 shows a schematic contract curve for the problem of carcinogen regulation.

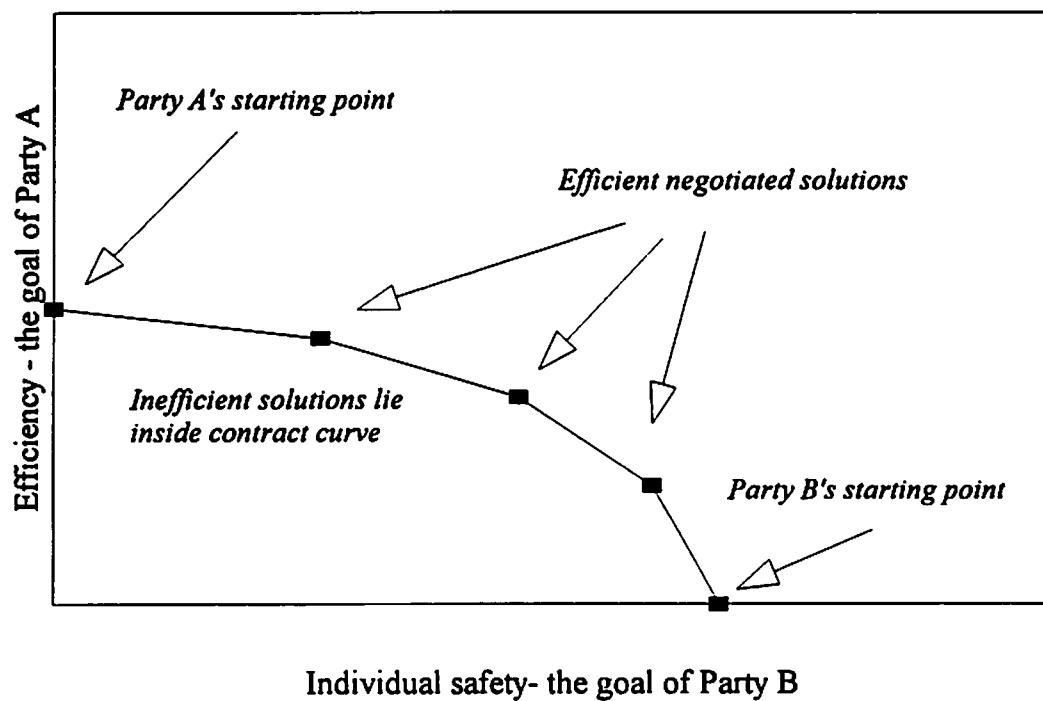


Figure 1. A schematic contract curve for carcinogen regulation framed as a negotiation problem.

The Positions of the Two Parties

Party A

In its role as the reviewer of regulations, the OMB is concerned about balance, wanting both the overall amount of regulation, and the set of regulations promulgated, to be optimal. Party A (the OMB) sets the problem in a neoclassical economic paradigm. The saving of statistical lives -- small reductions in small risks of death -- is a consumer good, like any other.¹⁸ In the absence of regulation, it can be expected that this good will be underprovided, since it is unpriced, and producers will have no incentive to provide it.¹⁹

Regulatory balance is to be sought by the use of cost-benefit analysis (CBA), and, for this, the benefit of averted cancer cases must be monetized. The methods developed within welfare economics for dealing with the general problem of valuing unpriced or underpriced goods are used by economists to address the problem of valuing statistical lives. These methods involve the inference of actual preferences for nonmarket goods by observing purchases of linked market goods (revealed preference), or the direct elicitation of preferences using surveys (contingent valuation). For example, revealed preferences for lifesaving have been studied by analyzing blue collar labor markets, since risky jobs may command a wage premium, and elicited preferences for lifesaving have been studied by

¹⁸ The first characterization of mortality risk reductions as a consumer good can be found in Schelling, 1968.

¹⁹ Carcinogen emissions are an external cost of production processes; carcinogen emission standards are intended to internalize this externality. In some situations, the cost of emissions may be internalized without regulation. Consider, for example, workers who are exposed to a risk, but are aware of it, willingly take it, and are compensated accordingly.

questioning drivers as to their willingness-to-pay for improvements in highway safety.²⁰

The values thus obtained for a statistical life range over an order of magnitude, with an upper limit at about ten million dollars.²¹ Statistical lives saved from cancer are a particular challenge for these methods. Almost all of the empirical work is directed at risks in which the cause of death is trauma from injury. Moreover, the requirement in CBA to discount benefits as well as costs leads to the choice of discount rate strongly influencing the value of statistical life for cancer deaths, because of the long latency period between exposure and diagnosis.²²

The OMB sidesteps the need to place a value on statistical life by evaluating regulations, not with CBA, but with a close methodological relative, cost-effectiveness analysis (CEA). A CEA is an economic analysis in which the result is reported as a ratio of cost per unit of effectiveness, thereby avoiding the monetization of effectiveness. Thus, carcinogen regulations can be compared with each other, and with other lifesaving interventions, using CE ratios of cost per life saved, or cost per life-year saved. Many toxin regulations save lives at costs well in excess of ten million dollars.²³

²⁰ Viscusi, 1993; Jones-Lee, 1989.

²¹ Fisher, *et al.*, 1989, Viscusi, 1993; Passell, 1994.

²² Discounting at a real rate of 3% to allow for a thirty year latency period would reduce the upper limit value placed on an averted cancer case from ten million dollars to about four million dollars.

²³ Tengs, *et al.*, 1995, found that the median cost per life-year (not cost per life) in 144 toxin control interventions was 2.8 million dollars.

Party B

In contrast with the OMB, the regulatory agencies (Party B) are driven by a variety of incentives to advocate the overprovision of risk reductions. The costs of risk reductions are borne by producers and consumers, not by the regulatory agencies. Moreover, in much of the legislation that governs the regulation of carcinogens, the regulatory agencies are forbidden to consider costs; their mandate is protection, not balance.²⁴

Thus, the focus of the regulatory agencies is on the benefits of risk reduction. Moreover, while the number of cases of cancer averted by a carcinogen regulation may be reported in a Regulatory Impact Analysis (RIA), the concept of an acceptable level of individual risk is key in the court rulings that constrain Party B's interpretation of the legislation.

Embedded in this concept of an acceptable level of risk is a switch -- a risk is either acceptable or not. This moral absolutism is very different from the utilitarian philosophy that underlies the neoclassical economic paradigm, in which values are set in context and everything is tradable.

In practice, the moral absolutism is weakened by pragmatism, because it is not always technically possible to reduce all risks to the level deemed acceptable. Hence, EPA's objective is a bright line of one in a million (*de minimis*), but one in ten thousand (*de*

²⁴ There are exceptions. The Federal Insecticide, Fungicide, and Rodenticide Act and the Toxic Substances Control Act require EPA to balance the benefits of proposed regulations against the costs; the Safe Drinking Water Act requires EPA to consider the costs of proposed rules, and not the benefits (Van Houtven & Cropper, 1996).

manifestis) will do.

In agency rule-making, some concern about cost is subsumed into “technical feasibility”, but only in terms of affordability for the industry as a whole, under the Regulatory Flexibility Act of 1980. This act is an attempt to protect jobs and industry viability. As part of an RIA for a significant rule, the agency must calculate compliance cost to revenue (or profit) ratios, with the intent of ensuring that the financial burden of a rule does not fall too heavily on a few firms or individuals. This is referred to as *economic feasibility*.

The formalization of Party B’s heuristic, necessary for placing the carcinogen regulation problem into a Raiffa-type negotiation model, is difficult. The negotiation framework is intended to resolve disputes over outcomes, whereas the regulatory agencies appear to be more concerned with process than outcome. This focus is a result of the constraining role played by the courts, since lawyers “*tend to look for solutions that are procedurally fair, or treat like persons alike, or that preserve past expectation more than they provide economically efficient solutions.*”²⁵

If the concept of an acceptable level of individual risk is to be taken literally, then those below a (*de manifestis*) bright line can be regarded as safe, and those above a bright line as in danger. Indeed, a prominent risk analyst has described a *de manifestis* risk as “*inherently unsafe*”.²⁶ Thus, the goal of Party B is taken to be to push as many people as possible below the bright line.

²⁵ Breyer, 1982, p.381.

²⁶ Travis, *et al.*, 1988, p.874.

The Negotiation Model -- Generation of a Contract Curve

The negotiation model is illustrated by working the problem through with a set of hypothetical potential carcinogen regulations. It would have been preferable to have used real data from Regulatory Impact Analyses (RIAs) prepared by agencies. However, this proved impossible, since agency analyses of proposed regulations do not provide all the data that one would like. A new approach requires new information. As described earlier in this paper, the first stage in a Raiffa-type negotiation analysis is the creative generation of many possible outcomes, and in preparing RIAs, regulatory agencies have already constrained the possibilities.²⁷

The objective is to get Party A, the advocate of allocative efficiency, and Party B, the advocate of individual safety, to agree on a set of differentiated standards, that is, a set of exposure limits that may differ across groups.²⁸

²⁷ In early work on this model, I experimented with real data from the 1994 OSHA cadmium rule, a carcinogen regulation that is representative of current OSHA policy. Two possible permissible exposure levels (PELs) for concentrations of cadmium dust in air for twenty industries were proposed in 1990. This appeared to give a good base for testing sets of differentiated standards. However, compliance costs were lower for the stricter PEL -- the opposite of what one would expect. This appears to be because more workers needed to wear respirators to achieve the stricter PEL, and respirators are cheaper than engineering controls. OSHA tries to minimize the use of respirators, because they are uncomfortable, and, thus, not as reliable as engineering controls. The use of this data would be confusing; the real world intrudes too much.

²⁸ Incentive-based mechanisms such as pollution taxes or tradable permits are, in theory at least, better ways to reduce carcinogen exposures than emission standards. But whatever the means, ends must still be chosen - the size of a Pigouvian tax or the amount of pollutant to be traded must be set. Some way of deciding which lives to save, when all cannot be saved, will always be needed.

Setting Up the Model

Different groups of people -- workers, residents, or consumers -- are exposed to various carcinogens as a result of industrial production. Table 1 shows the hypothetical data used to illustrate the model.²⁹

For simplicity, a single carcinogen is considered, but the method is generalizable to the simultaneous regulation of many carcinogens. In the current regulatory process, standards are set for one carcinogen at a time.³⁰ Reducing the risk from a single carcinogen to an acceptable level has little ethical significance, since different groups are exposed to different numbers of carcinogens. For example, workers in one industry may be exposed to a single carcinogen, and workers in another industry exposed to that carcinogen along with twenty others. An acceptable level of individual risk would gain much more ethical significance, if all carcinogens of industrial origin were to be regulated in one exercise. This is an impossible ideal. However, it may be possible to set standards for several carcinogens together, especially if they threaten the same groups.

Table 1 shows ten groups of people (a through j) exposed to the carcinogen. Within

²⁹ In creating this data, I have assumed that dominated strategies have been discarded. Within each group, the incremental cost of saving a life, and the incremental cost of pushing an individual below the bright line, both rise as the stringency of the standard increases.

³⁰ EPA standards and policies are actually developed separately for the same carcinogen from *different sources*. For instance, radon emissions into outdoor air from industrial and military sources are regulated under the Clean Air Act, radon concentrations in residential water systems are regulated under the Safe Drinking Water Act, and recommendations for radon concentrations in indoor air are developed under the Indoor Radon Abatement Act. The results -- in terms of limits on individual risk -- are, not surprisingly, very different.

each group, the risk will vary. Distributions of individual risk from pollutants are typically modelled as lognormal, with many exposed to small risks, and few exposed to large risks.³¹

There are two ways of protecting a group with a given risk distribution. Either engineering controls can be used that will shift the whole distribution to the left, or the high end of the distribution can be truncated by reducing the exposure of the very high risk individuals. In general, the latter approach relies on behavioral changes, such as requiring workers to wear respirators and protective clothing, or reducing the consumption of certain fruits and vegetables containing carcinogenic residues by small children; this approach has not been favored by the regulatory agencies. The risk reduction in the model is achieved by the use of engineering control technologies, and these are assumed to be “lumpy”. For each group exposed to the carcinogen, there are four potential levels of protection -- the status quo, a slack standard, a moderate standard, and a tight standard.³² The code a1 denotes the slack standard for the first exposed group, and so on.

For each group, for each level of protection, the annual number of lives lost is given, and the annual number of lives saved by each level of protection can be calculated by simple subtraction.³³ The same applies to the number of people in each group who are not “safe”,

³¹ The risk in this analysis are all mortality risks. It is assumed that the carcinogens have no other deleterious effects on health.

³² No judgement of the worth of the standards is intended by the use of the adjectives “slack”, “moderate”, and “tight”.

³³ The issue of discounting has been set aside in this illustration of the model, since whether the saved lives are discounted or not makes no difference to the negotiation analysis. Discounting would make a difference if more than one carcinogen was being assessed, and the carcinogens had different latency periods, and/or were diagnosed at different mean ages.

because their risk level is above the bright line.

For each group, for each standard, the annual compliance cost is given. The costs of the different standards should ideally be the sum of the loss of producer surplus and the loss of consumer surplus, since regulation leads to reduced profits for industry and higher prices for consumer products. However, such estimations of costs are usually an ideal rather than a reality.

Table 1. Hypothetical data used to illustrate the negotiation model.

Status quo

Group code	Annual lives lost	People above bright line
a	1.14	5,000
b	1.15	10,000
c	2.74	18,000
d	1.36	35,000
e	1.83	38,000
f	5.50	17,000
g	4.41	27,000
h	1.41	20,000
i	1.85	25,000
j	0.94	5,000
	22.33	200,000

Slack standard

Moderate standard

Tight standard

Code	Annual cost M\$	Annual lives lost	People above bright line	Code	Annual cost M\$	Annual lives lost	People above bright line	Code	Annual cost M\$	Annual lives lost	People above bright line
a1	8.4	0.57	2,000	a2	7.0	0.20	1,200	a3	5.4	0.13	300
b1	0.6	0.78	8,000	b2	2.3	0.47	5,000	b3	2.1	0.26	4,000
c1	4.7	1.35	11,000	c2	10.0	0.79	8,000	c3	10.6	0.56	5,000
d1	18.8	0.63	12,000	d2	10.0	0.44	3,000	d3	4.1	0.39	200
e1	12.2	1.11	20,000	e2	10.0	0.94	10,000	e3	6.9	0.84	9,000
f1	21.0	2.22	11,000	f2	11.8	1.43	8,000	f3	11.3	0.81	6,000
g1	13.9	1.93	10,000	g2	10.0	1.58	7,000	g3	5.3	1.46	6,000
h1	7.0	0.62	7,000	h2	6.8	0.33	2,000	h3	4.5	0.25	700
i1	11.6	0.47	4,000	i2	2.2	0.24	1,000	i3	1.7	0.13	500
j1	8.3	0.73	1,000	j2	6.6	0.64	300	j3	2.8	0.61	200
	106.5	10.40	86,000		76.6	7.06	45,500		54.8	5.44	31,900

Each Party Ranks the Differentiated Standards

Party A wishes to perform a cost-benefit analysis and must, therefore, monetize the averted cancer cases in order to calculate benefits. The difficulties in doing this have been discussed earlier in this paper, and a threshold value of ten million dollars for an undiscounted saved life will serve the pedagogical purpose here.

Recall there are three possible standards for each group. Party A balances costs and benefits on the margin, so for each group, the incremental costs and benefits of each of the three standards is calculated -- the slack standard compared with no standard, the moderate standard compared with the slack standard, and the tight standard compared with the moderate standard. Net benefit is positive whenever the cost of saving a life is less than ten million dollars, and Party A finds these cases by ordering the standards by increasing cost per life saved.³⁴

Party B wishes to push everyone below the bright line of acceptable risk, and, unless checked in some way, will go to the limit of technical feasibility.

The ranking of the differentiated standards by both parties is shown in Table 2. Were a regulatory budget to be imposed, each party would work down its ordering until the budget was exhausted. In the lower half of Table 2, the effect of setting an annual regulatory budget of 100 million dollars is shown. In this illustration, Party A would not be constrained by the budget, but would choose only those standards, for which the cost of saving a life is less than

³⁴ Party A could rank potential standards, not by increasing incremental cost of saving a *life*, but by increasing incremental cost of saving a *life-year*. This small change would allow for recognition of the age of death.

ten million dollars. Party B would be constrained by the budget. Note, however, that the parties agree on the standard that should be set for five of the groups. Although Party B generally favors tighter standards than Party A, the opposite is the case for two groups, showing that Party B passes up the opportunity to save some cheap lives.

Table 2. Ranking of the differentiated standards by the two parties.

Party A					Party B				
Code	Cost M\$/yr	Annual lives saved	Cost per life M\$/life	Total cost M\$/yr	Code	Cost M\$/yr	People made "safe"	Cost per person made "safe" (\$)	Total cost M\$/yr
b1	0.6	0.37	1.5	1	b1	0.6	2,000	277	1
c1	4.7	1.39	3.4	5	h1	7.0	13,000	539	8
g1	13.9	2.48	5.6	19	i1	11.6	21,000	552	19
f1	21.0	3.28	6.4	40	c1	4.7	7,000	675	24
b2	2.3	0.31	7.3	42	e1	12.2	18,000	676	36
i1	11.6	1.38	8.4	54	i2	2.2	3,000	744	38
h1	7.0	0.79	8.9	61	b2	2.3	3,000	754	41
i2	2.2	0.23	9.7	63	g1	13.9	17,000	817	54
b3	2.1	0.21	9.8	65	d1	18.8	23,000	819	73
a1	8.4	0.58	14.6	74	e2	10.0	10,000	1,000	83
f2	11.8	0.79	14.9	85	d2	10.0	9,000	1,110	93
i3	1.7	0.11	15.8	87	h2	6.8	5,000	1,363	100
e1	12.2	0.72	16.9	99	d3	4.1	2,800	1,461	104
c2	10.0	0.56	17.9	109	b3	2.1	1,000	2,058	106
f3	11.3	0.62	18.3	121	j1	8.3	4,000	2,084	115
a2	7.0	0.37	19.0	128	a1	8.4	3,000	2,800	123
h2	6.8	0.29	23.5	135	g2	10.0	3,000	3,337	133
d1	18.8	0.73	25.8	153	c2	10.0	3,000	3,341	143
g2	10.0	0.35	28.6	163	h3	4.5	1,300	3,434	147
j1	8.3	0.21	39.7	172	i3	1.7	500	3,476	149
g3	5.3	0.12	44.3	177	f1	21.0	6,000	3,499	170
c3	10.6	0.23	46.2	188	c3	10.6	3,000	3,542	181
d2	10.0	0.19	52.6	198	f2	11.8	3,000	3,924	193
h3	4.5	0.08	55.8	202	g3	5.3	1,000	5,316	198
e2	10.0	0.17	58.8	212	f3	11.3	2,000	5,673	209
e3	6.9	0.10	69.4	219	a3	5.4	900	6,012	215
j2	6.6	0.09	72.8	226	e3	6.9	1,000	6,940	222
a3	5.4	0.07	77.3	231	a2	7.0	800	8,690	229
d3	4.1	0.05	81.8	235	j2	6.6	700	9,360	235
j3	2.8	0.03	94.3	238	j3	2.8	100	28,290	238

Suppose an annual regulatory budget is set at 100 million dollars:

Party A is unconstrained.
Standards chosen by Party A

a Status quo
b Tight
c Slack
d Status quo
e Status quo
f Slack
g Slack
h Slack
i Moderate
j Status quo

Party B is constrained.
Standards chosen by Party B

a Status quo
b Moderate
c Slack
d Moderate
e Moderate
f Status quo
g Slack
h Moderate
i Moderate
j Status quo

Generation of the Contract Curve

Now that the decision heuristics of both parties have been formalized into ranking algorithms, the contract curve can be generated. It is assumed that a regulatory budget has been set, limiting the total amount of money that producers can be compelled to spend to reduce exposure to this carcinogen.³⁵ Suppose that the budget is set at the total cost of the slack standard, uniformly applied across all ten groups, that is, 106.5 million dollars per year.

Both parties perform a simple constrained optimization to find their starting positions for the negotiation, represented by the extreme points on the contract curve. The intermediate points on the contract curve are found by maximizing a combined objective function --

$$\{ \alpha * \sum \text{Lives saved} + \beta * \sum \text{Individuals made "safe"} \}$$

-- using different values of α and β to vary the weighting on each of the component objectives.³⁶

Table 3 shows a sample optimization with α set equal to one and β set equal to zero, which corresponds to Party A's initial position. Total annual cost is constrained to be 106.5 million dollars, and the solutions are constrained to be non-negative, and not greater than unity. A solution of unity denotes acceptance of the standard, and a solution of zero denotes

³⁵ Producers pay the direct costs of regulations, but share the burden of those costs. Some fraction of the compliance cost of a regulation is passed on to consumers. Producers may choose to pay for a regulation by reducing the wages of the workers protected by the regulation.

³⁶ Only one weighting factor -- α or β -- is actually required, but two are presented here in the spirit of explicit recognition of the interests of both parties.

rejection of the standard. The solutions are not constrained to be integers, since this gives a “cleaner” result, so that the standard on the margin between acceptance and rejection has a fractional solution.

The contract curve is shown in Figure 2, tracing out the opportunity cost associated with each party’s choices, in terms of the other party’s objective. The two parties have a great deal in common. Note that the uniform slack standard, which saves 11.93 lives per year, and pushes 114,000 individuals below the bright line, lies within the contract curve.

This illustration has shown that the conflict between the OMB and the regulatory agencies over the choice of carcinogen standards can be analyzed in the analytic framework for negotiation developed by Raiffa. The two parties may have a great deal in common, and uniform standards are optimal for neither. However, the ranking algorithm created by the formalization of Party B’s decision heuristic is problematic for two reasons.

First, the simple application of negotiation analysis used above assumes that only *outcomes* are valued, whereas Party B may value the current *process* of carcinogen regulation very highly. This does not seem likely; Breyer describes the process of risk regulation as so arduous and adversarial, that agencies may choose to promulgate rules that they know could be improved, but fear another round of judicial review.³⁷

Second, the goal of pushing as many people as possible below the bright line is hard to take seriously. Is this really what Party B would try to do if constrained by a regulatory budget? The formalization of the regulatory agencies’ heuristic has led to a ranking algorithm that brings into question the heuristic itself.

³⁷ Breyer, 1982, p.109.

Table 3. A sample constrained optimization for mapping the contract curve.

The objective function is $\{\alpha * \text{lives saved} + \beta * \text{people made "safe"}\}$

The regulatory budget is set at 106.5 million dollars per year, the cost of a uniform slack standard.

This set of solutions is optimal for Party A.

Standard code	Cost M\$/yr	Lives saved per yr	People made "safe"	Solutions
a1	8.4	0.58	3,000	1
b1	0.6	0.37	2,000	1
c1	4.7	1.39	7,000	1
d1	18.8	0.73	23,000	0
e1	12.2	0.72	18,000	1
f1	21.0	3.28	6,000	1
g1	13.9	2.48	17,000	1
h1	7.0	0.79	13,000	1
i1	11.6	1.38	21,000	1
j1	8.3	0.21	4,000	0
a2	7.0	0.37	800	0
b2	2.3	0.31	3,000	1
c2	10.0	0.56	3,000	0.71
d2	10.0	0.19	9,000	0
e2	10.0	0.17	10,000	0
f2	11.8	0.79	3,000	1
g2	10.0	0.35	3,000	0
h2	6.8	0.29	5,000	0
i2	2.2	0.23	3,000	1
j2	6.6	0.09	700	0
a3	5.4	0.07	900	0
b3	2.1	0.21	1,000	1
c3	10.6	0.23	3,000	0
d3	4.1	0.05	2,800	0
e3	6.9	0.10	1,000	0
f3	11.3	0.62	2,000	0
g3	5.3	0.12	1,000	0
h3	4.5	0.08	1,300	0
i3	1.7	0.11	500	1
j3	2.8	0.03	100	0
Objective		13.03		
Budget (M\$)		106.50		
Alpha		1		
Beta		0		
Total annual lives saved		13.03		
Total people made "safe"		99,628		




Figure 2. A contract curve for the imaginary carcinogen regulation problem developed in the text. The tradeoff between the objectives of the two parties is shown explicitly. The regulatory budget is set at 106.5 million dollars per year, the compliance cost of the "slack" standard applied uniformly across all exposed groups.

A Critique of the Bright Line Heuristic .

The concept of an acceptable level of individual risk (a *de manifestis* bright line) has logically led to a ranking algorithm for the allocation of a regulatory budget that yields transparently unfair outcomes. The gain from reducing one person's risk from a high *ex ante* level to the bright line is given the same weight as the gain from reducing another person's risk from a low *ex ante* level to the bright line. The value placed (implicitly) on the first person's life is less than the value placed on the second person's life. Although the goals of efficiency and equity are popularly viewed to be at loggerheads, it can be argued that Party A's heuristic leads to fairer outcomes than Party B's heuristic, since Party A places the same value on all lives.

The evolution of the bright lines used by FDA and EPA has already been described. There is clearly an anchoring on levels of individual risk rounded to orders of magnitude -- one in a million, one in ten thousand -- this alone is evidence for the arbitrariness of the selection of bright lines. OSHA's *de manifestis* bright line of one in a thousand is based on a much-quoted suggestion in an opinion written by a Supreme Court Justice in 1980.

*"It is the Agency's responsibility to determine, in the first instance, what it considers to be a 'significant' risk. Some risks are plainly acceptable, and others are plainly unacceptable. If, for example, the odds are one in a billion that a person will die from cancer by taking a drink of chlorinated water, the risk clearly could not be considered significant. On the other hand, if the odds are one in a thousand that regular inhalation of gasoline vapors that are 2% benzene will be fatal, a reasonable person might well consider the risk significant and take appropriate steps to decrease or eliminate it."*³⁸

³⁸ Justice John Paul Stevens in the 1980 benzene case (Industrial Union Department, AFL-CIO v. American Petroleum Institute, 448 U.S. 607, 1980, p.655).

This suggestion has been adopted as a policy, but it has no underlying rationale with a basis in theory. In contrast, the decision heuristic of allocative efficiency used by Party A - the OMB - does have a basis in the theory of welfare economics.

The rationales for numerical bright lines that have been proffered appear to be based on the premise that consistency *per se* is desirable. Patterns are sought in past decisions, and the argument made that future decisions should be consistent with these revealed regulatory preferences. Two of these rationales are now examined.

Milvy (1986) proposes a guideline for managing risk from carcinogens using a three part argument. The first part is that the acceptable level of individual risk should depend on the size of the population at risk -- deaths are more acceptable if they are perceived as more random, that is, the value of a life rises with its *ex ante* identifiability. In the second part of his argument, Milvy anchors acceptable risk on two benchmarks - one in a million as an acceptable level for the whole U.S. population (because a "*consensus*" on this exists), and 2.1 in a hundred thousand as an acceptable level for a hundred people, because it is close to the lifetime risk of a fatal accident on the job borne by white collar workers. The former is used to determine an acceptable population risk or body count; it is apparently acceptable for one out of every million Americans to die from a particular industrial carcinogen. The latter is used to determine an acceptable maximum individual lifetime risk from the same carcinogen. From these, Milvy creates an equation in which the acceptable level of individual risk varies with the inverse square root of the size of the affected population. Finally, he fits a polynomial to risk and population data from a set of regulatory decisions (with no statistic to indicate "the goodness of fit"), in which the relationship between acceptable risk and

population varies in the same way as that in his equation derived from the two benchmarks. The result is a kind of middle path, balancing acceptable population risk against acceptable individual risk.

This guideline has been presented in some detail, since it has been taken seriously by the academic community. However, it is simply an argument from history -- the best thing to do in the future is to continue what has been done in the past -- and its complexity serves to protect it from scrutiny.³⁹ No reason is given for why population risk should matter, yet from an economic perspective, this is the strongest part of the argument. A low body count *ex ante* means that the monetized benefit of the regulation must be low, and the transaction costs alone may well exceed the benefit.

A much clearer empirical and prescriptive study of carcinogen management has been done by Travis, *et al.* (1987a, 1987b, 1988). The authors examined 132 regulatory decisions made by several federal agencies, and found a surprising consistency. Like Milvy's proposal, their prescription for carcinogen management is an argument from history, because it is based on the assumption that past decisions were correct.

Travis, *et al.* search for bright lines of individual risk implicit in past decisions -- the *de manifestis* risk, that is "*inherently unsafe*"⁴⁰ and must be reduced, and the *de minimis* risk, that is so small it should be ignored. They argue that federal agencies have always acted to reduce risks when the lifetime risk to an individual is above four in a thousand - this is the

³⁹ A similar, more recent attempt by Finkel, 1990a, does not make Milvy's proposal any more convincing.

⁴⁰ Travis, *et al.*, 1988, p.874.

historic *de manifestis* level. The historic *de minimis* level is the familiar one in a million. For the relatively few cases (23 of the 132) in which the individual risk lay between the *de manifestis* and *de minimis* levels, costs have apparently been balanced against benefits, with the threshold value of life at about two million dollars. The authors argue that these revealed decision rules are acceptable as a policy prescription, provided the *de manifestis* level is allowed to vary with the size of the population at risk.⁴¹

Levels of individual risk above the *de manifestis* level are then *unacceptable*, implying that the cost of their reduction is irrelevant. It can be argued that it is reasonable for OSHA to adopt a permissive posture, because the cost of carcinogen emissions is linked to the benefit of employment for the exposed workers, and a portion of their wages may be viewed as a risk premium. FDA and EPA should have stricter bright lines, since those who bear the risk do not benefit from the industrial processes that generate the risk any more than the average person. However, the choice of the actual levels of acceptable risk remains entirely arbitrary; they are benchmarked to past decisions, and rounded to orders of magnitude.

There is one rationale for the notion of an acceptable level of risk that may be strongly rooted in theory -- biological theory. If the dose-response relationship for a carcinogen has a threshold, then doses below the threshold are indeed inherently safe, and the *de manifestis* level could be set at the threshold.

Little is known about human dose-response relationships at low doses, and for most

⁴¹ A *de manifestis* level of one in a thousand is "*tolerable if reduction of exposures would not appreciably reduce total population cancer incidence*" (Travis & Hattemer-Frey, 1988, p.876). This is again the argument of balancing acceptable population risk against acceptable individual risk.

suspected carcinogens, dose-response relationships are extrapolations from experiments with rodents. Because such experiments cannot involve millions of rodents, the doses must be much larger than those typical of human exposures, in order that differences between test and control groups can be observed. Some scientists expect that many carcinogen dose-response relationships are shaped like hockey sticks; that is, there is a threshold below which a dose elicits no response. Most scientists believe that genotoxic carcinogens, which have the ability to interact directly with DNA, are unlikely to have thresholds. If acceptable levels of individual risk were to be based on threshold doses, they would have to be specific to the carcinogen.⁴² However, progress in determining the functional form of dose-response relationships, at the low doses typical of most carcinogen exposures, is very slow, and this approach cannot be used without a quantum leap in scientific understanding.

Even without formalization into a ranking algorithm, the bright line heuristic makes little sense.⁴³ However, the notion of an acceptable level of *individual* risk focuses attention on the *distribution* of risk. If the risk from a carcinogen is distributed unevenly across individuals, we may want to give some priority to those who bear the heaviest risk burdens. Party B needs a better decision heuristic, one that has a sound rationale, and allows weighting for *degrees of unacceptability* of individual risk.

⁴² Note that *inherent safety* would continue to be elusive, since the reduction of one risk is generally associated with increases in other risks. For example, reduction of carcinogenic residues from pesticides in food may drive up the prices of fruits and vegetables, lowering consumption of some that reduce cancer risk. See Graham & Wiener, 1995.

⁴³ In reality, bright lines of individual risk are probably viewed by their advocates as a strategy for “holding the line” against industry, not as boundaries between safety and danger.

Improving the Heuristic of the Regulatory Agencies

An improved decision heuristic (and ranking algorithm) may be found for Party B in the prescriptive “value-of-life” models developed by decision analysts working on the general problem of lifesaving choices.⁴⁴ A number of scholars have explored the properties of multiattribute utility functions in which changes in consumption are traded against changes in risk.⁴⁵ The underlying assumption is that reductions in risks of death are worth what people are prepared to pay for them.

This is, of course, the same rationale – willingness-to-pay (WTP) -- used by Party A, as the basis of its heuristic. However, these models show that the same value of life should *not* be used for all lifesaving decisions, and that it should vary with the magnitude of the risk.

⁴⁴ An alternative place to look for a decision heuristic that captures a concern for fairness is the welfare economics literature on income distribution. One candidate is an adaptation of Sen’s poverty index. Sen (1976) reviewed a number of measures of income inequality and proposed a new measure of poverty that avoids the shortcomings of those then in use, namely:

$$P = H [I + (1-I) G]$$

H = the percentage of the population below the poverty level, the *head-count ratio*.

I = the percentage of the mean short-fall from the poverty level, the *income-gap ratio*.

G = the *Gini coefficient* of the income distribution of the poor.

Finkel (1990b) has suggested adapting the Sen index for use in carcinogen regulation, using acceptable risk as the equivalent of the poverty level, but this will not do. The poverty level is a benchmark of something real -- shelter, enough to eat, and so on -- in the sense of Rawlsian primary goods, whereas bright lines of acceptable risk are entirely arbitrary. In order for a bright line of acceptable risk to play a role analogous to the poverty level, it would have to represent a level of inherent safety, which it does not.

⁴⁵ Linnerooth (1979) reviews the early models of this genre. The model developed by Howard (1980) is the focus in the second half of this paper. Jones-Lee has been a major contributor (e.g. Jones-Lee, 1989). For a recent example, see Pratt & Zeckhauser, 1996.

Hence, a model of this type might provide a scheme for weighting the value of life according to the distribution of the risk. Such a weighting scheme would allow Party B to modify Party A's ranking algorithm into a form that would explicitly recognize the distribution of the risk over individuals. *Since Party A's concern for efficiency and Party B's concern for the protection of individuals would both be captured in the modified algorithm, both parties should agree on the allocation of the regulatory budget, and there would be no need for a negotiation.*

Hammitt (1990) presents the general result obtained from the value-of-life models. An adaptation is shown in Figure 3. The price an individual is willing to pay for unit reductions in mortality risks can be expected to vary with the size of the risk reductions. As risk reductions become larger, the individual's baseline risk is increasingly affected, and WTP for a unit of risk reduction rises ever more steeply. At any point on the curve, the implicit value of life is the slope.

This general result fits well with our intuitions about how we would expect the value of life to vary with the size of a risk. Consider Figure 3 to be composed of three regions corresponding to three risk levels - *very small*, *very large*, and *in-between*.




Figure 3. A general model showing how willingness-to-pay can be expected to vary for risk reductions, adapted from Hammitt, 1990, p. 134. By way of illustration, the value of statistical life is taken as eight million dollars; the numbers on the axes should not be taken literally.

First, over a range of *very small* mortality risks, it is reasonable to assume that WTP for a unit risk change is constant, since none of the risks is large enough to substantially affect the individual's baseline risk. The WTP function is linear in this small-risk region. The slope, and, therefore, the implicit value of life is constant and is known in the literature as, *the value of statistical life*. I follow Howard (1980) in using the term *small-risk value of life*.

Second, for *very large* mortality risks, the WTP curve rises steeply. At some risk level below the probability of certain death ($p = 1$), the WTP curve rises asymptotically to infinity. The slope, the implicit value of life, behaves in the same way. When facing very high risks of death, an individual will pay a great deal to improve the odds. But, at some point, the individual's income will begin to exert a constraining effect on WTP. Well before the function becomes vertical, the individual simply could not afford to purchase risk reductions.

The relevant question for the individual involuntarily exposed to a carcinogen of industrial origin is *not* how much he or she is willing to pay for the reduction of that risk, but how much he or she is *willing to accept as compensation* for bearing that risk. WTA diverges increasingly from WTP as the risk rises, since the latter is constrained by the individual's income, and the former is not. The asymptotic rise to infinity of the function in Figure 3 only makes sense if the function is WTA, not WTP. At some level of risk, short of certain death, no amount of money could compensate the individual for bearing that risk, and both the compensation required, and the value of life, would be infinite.

The decision scientist, Ron Howard, has dealt with the WTP - WTA distinction very clearly in a remarkable value-of-life model.⁴⁶ Howard sets up two problems. In the White

⁴⁶ Howard, 1980; 1984.

Pill problem, the individual buys a decreased risk of death (WTP); in the Black Pill problem, the individual is compensated for bearing an increased risk of death (WTA). Carcinogen regulation is a Black Pill problem, since the risk is borne involuntarily, and, in the next section of this paper, a simple variant of Howard's model is developed.

Finally, for *in-between* risks, the (now) WTA function in Figure 3 rises at an increasing rate, as the risk moves from the small-risk region to the very high risk region. The slope, the implicit value of life, behaves similarly.

Thus, an individual seeking monetary compensation for involuntarily imposed risks would logically weight his or her value of life in accordance with the *slope* of the function in Figure 3. The value of life should be constant over the small-risk region, and then rise at an increasing rate, finally approaching infinity at some risk that is less than certain death. If such a weighting schedule can be developed, then the basis exists for a new ranking algorithm for the carcinogen regulation problem.

A model is required whose solution will be a function, $V = V(p)$, where V is the implicit value of life, and p is the lifetime risk to an individual, from a single carcinogen, or set of carcinogens. A small-risk value of life, V_s , can be taken from the empirical studies on revealed and elicited preferences for reductions in risk. V_s will be in the range of two to ten million dollars.

Suppose a proposed carcinogen standard reduces the lifetime risk of each of N individuals from p_1 to p_2 . Party A would measure the benefit of this standard as:

$N * (p_1 - p_2)$ saved lives and monetize the benefit as:

$$N * (p_1 - p_2) * V_s$$

If, however, the value of life is allowed to vary with the distribution of the risk burden, then the benefit of the proposed standard can be monetized as:

$$N * \{ p_1 * V(p_1) - p_2 * V(p_2) \}$$

There are two major tasks remaining. First, the function, $V = V(p)$ ⁴⁷ must be constructed for the individual who must balance compensatory payments against increased risks of death. Second, the result for the individual decision-maker must be extended to a result for a group of individuals, where there is a distribution of risk over the individuals in the group; this distribution will shift if a carcinogen standard is set.

⁴⁷ Individual lifetime risk (p) is, in several ways, a poor metric, although it serves the purpose for the analysis that follows. Its replacement by loss of life expectancy (LLE), measured in days, would have both substantive and cognitive merit. Substantively, one could distinguish between a carcinogen that generally kills people in their fifties and another that generally kills people in their seventies. There would be several cognitive advantages. The use of LLE would provide perspective on the benefits of carcinogen regulation, allowing them to be benchmarked to the benefits of other lifesaving investments, such as cancer screening tests. People exaggerate small probabilities, and bright lines “sound” bigger than they are. Assuming that the average loss of life for someone who dies from cancer because of a critical exposure to a carcinogen is 16 years, then a bright line of one in a thousand corresponds to six days of LLE, and a bright line of one in ten thousand to 14 hours of LLE. Risks expressed as ratios with very large denominators are confusing; one in five hundred sounds smaller than one in a thousand to many people. Finally, risk distributions for carcinogens are usually plotted on a semi-logarithmic scale; it is easy to forget, when studying such distributions, that unit changes in scale are, in fact, order of magnitude changes.

The Black Pill Model

In this section, a weighting function for the value of life, $V = V(p)$ is developed from a simplification of Howard's Black Pill model.⁴⁸

Consider an individual, with a lifetime consumption of W , offered monetary compensation of a lump sum of x dollars, for accepting a black pill, that is, an additional probability of death, p . The choice is depicted in Figure 4 as a decision tree.

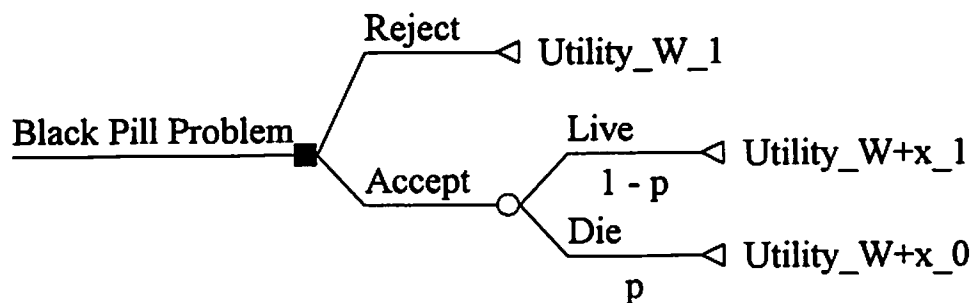


Figure 4. The Black Pill decision problem. The individual must decide whether or not to accept the black pill, a risk of death, p . The individual's initial lifetime consumption is denoted W . If the black pill is accepted, monetary compensation of x dollars is paid.

⁴⁸ The solution of Howard's more complete model has the same features as that of the simple one-period variant presented here. There is no advantage to adopting the complexity of his model, since it does not "fit" the carcinogen regulation problem well. Although Howard models the remaining lifetime of the individual facing the Black Pill decision year by year, there is no simple way to incorporate the latency period between exposure to the carcinogen and diagnosis of a (fatal) tumor.

The individual's utility function has two attributes, consumption W , and life, Λ .

$$U(W, \Lambda) = U(W, 1) \text{ if the individual lives.}$$

$$U(W, \Lambda) = U(W, 0) \text{ if the individual dies.}$$

The individual is indifferent between rejecting and accepting the black pill when:

$$U(W, 1) = (1 - p) * U(W+x, 1) + p * U(W, 0) \quad (1)$$

It is assumed that:

$$U(W, 0) = U(0, 1)$$

implying that zero consumption is tantamount to being dead; and that:

$$U(W, 0) = U(0, 0)$$

implying that there is no utility from bequests.

Then:

$$U(W, \Lambda) = U(W)$$

and equation (1) becomes:

$$U(W) = (1 - p) * U(W+x) + p * U(0) \quad (2)$$

which can be rearranged to give an expression for p .

$$p = [U(W+x) - U(W)] / [U(W+x) - U(0)] \quad (3)$$

Now, again following Howard, *an illustrative example* can be constructed.

Suppose the following:

- The individual is risk-averse with respect to consumption, and the coefficient of risk aversion is constant. A simple exponential function has these properties, and is analytically tractable if constrained between -1 and 0.

$$U(W) = -\exp\{-W/\rho\} \text{ where:}$$

ρ is the risk tolerance, defined as the inverse of the coefficient of risk aversion.

- Initial consumption is 600,000 dollars. Since the model has only one period, this represents lifetime consumption. It is assumed that the average individual wishes to consume what he or she produces. Annual consumption throughout the individual's lifetime is taken to be constant, and equal to the 1995 GDP per capita.⁴⁹
- The risk tolerance, ρ , is 180,000 dollars. This implies that the individual is indifferent between a certain lifetime consumption of 600,000 dollars and a gamble, in which the chance of winning and receiving a high consumption of 790,000 dollars is roughly equal to the chance of losing and receiving a low consumption of 510,000 dollars.⁵⁰

For the illustrative example, equation (2) becomes:

$$p = [-\exp\{-(W+x)/\rho\} + \exp\{-W/\rho\}] / [-\exp\{-(W+x)/\rho\} + 1]$$

and the size of the black pill, p , can be calculated for any compensatory payment, x .

⁴⁹ The discount rate used is 5%. The results of this analysis do not hinge on an accurate estimate of lifetime consumption.

⁵⁰ The expected value of this gamble is \$650,000. An individual who is risk-averse with respect to wealth always prefers a certain outcome to a gamble with the same expected outcome.

Since, by definition, the value of life $V(p)$ is:

$$V = x / p \quad (4)$$

V can be mapped out as a function of p .

Table 4 shows the calculations for the base case of the illustrative example. The value of life V (actually, its logarithm to base 10) is mapped out against p in Figure 5. It can be seen that, indeed, this plot of V against p follows the general pattern of the slope of the function in Figure 3.

Table 4. The implicit value of life calculated using the simple black pill model.

The small-risk value of life is 4.87 million dollars.

The maximum probability of death is 3.57%.

Wealth	600,000		
Risk tolerance	180,000		
WTA compensation payments x (\$)	Black pill % probability of death p	log(V)	Implicit value of life (million \$) V
1	0.00002	6.69	4.87
10	0.0002	6.69	4.87
100	0.002	6.69	4.87
1,000	0.02	6.69	4.88
5,000	0.10	6.69	4.94
10,000	0.20	6.70	5.01
20,000	0.39	6.71	5.16
40,000	0.73	6.74	5.47
60,000	1.04	6.76	5.78
80,000	1.31	6.79	6.11
100,000	1.55	6.81	6.44
120,000	1.77	6.83	6.79
140,000	1.96	6.85	7.14
160,000	2.13	6.88	7.50
180,000	2.29	6.90	7.88
200,000	2.42	6.92	8.26
250,000	2.70	6.97	9.25
300,000	2.91	7.01	10.30
400,000	3.19	7.10	12.53
500,000	3.35	7.17	14.91
600,000	3.44	7.24	17.42
700,000	3.50	7.30	20.02
800,000	3.53	7.36	22.68
900,000	3.54	7.40	25.39
1,000,000	3.55	7.45	28.14
2,000,000	3.57	7.75	56.06
5,000,000	3.57	8.15	140.16
10,000,000	3.57	8.45	280.32
50,000,000	3.57	9.15	1401.58
100,000,000	3.57	9.45	2803.16
1,000,000,000	3.57	10.45	28031.62




Figure 5. The base case result for the illustrative example of the simple Black Pill model. The order of magnitude of the implicit value of life is shown as a function of the size of the risk.

In the small-risk region, the value of life is constant at just under five million dollars. At a probability of death far below 100%, the value of life rises rapidly to infinity; there is a limit on the size of the black pill that the individual can be “bribed” to accept. The limiting probability is now denoted p_{\max} , and, in this case, is equal to 3.57%.⁵¹ For risks that are in the middle “in-between” region – neither so small that weighting the value of life is not worth the trouble, nor so large that they will never be accepted -- $V(p)$ rises exponentially.

The two key results of a constant small-risk value of life (V_s), and a p_{\max} far below 100% are always obtained when a simple exponential utility function is used, regardless of the values taken by the parameters. A greater tolerance for risk is reflected in a lower small-risk value of life and a higher limiting probability; this is shown in Figure 6. A greater initial consumption is reflected in a higher small-risk value of life and a lower limiting probability; this is shown in Figure 7.

⁵¹ By way of comparison, Howard (1980), using the same utility function, and the same initial utility, found a small-risk value of life equal to 2.43 million dollars (about five million dollars if adjusted for inflation), but a higher p_{\max} of about 10%. The difference is due to his use of a many-period model, and the consequent discounting of future utility.




Figure 6. Variation in the result of the illustrative example of the simple Black Pill model. The sensitivity of the small-risk value of life, V_s , and the limiting probability, p_{max} , to changes in the individual's risk tolerance with respect to consumption is shown.




Figure 7. Variation in the result of the illustrative example of the simple Black Pill model. The sensitivity of the small-risk value of life, V_s , and the limiting probability, p_{max} , to changes in the individual's lifetime consumption is shown.

Extension of the Result to a Group and The Policy Implication

The model must now be extended to change the perspective from that of an individual decision-maker to that of an agent of society making decisions about the protection of groups of individuals. The first step in making this transition is to ensure that the underlying rationale, the basis on which the decision is to be made, is transferrable from the individual perspective to the social perspective.

The individual chooses whether or not to accept the black pill by considering the compensation that is offered. As the black pill risk becomes larger, the compensation demanded accelerates. The underlying principle is that of *fair compensation for voluntarily borne risks*.

Industries that generate carcinogenic pollutants as an external production cost could be required to *actually* compensate those who bear the risks. That possibility is outside this analysis. However, the Kaldor-Hicks criterion, that compensation need only be *potential*, not *actual*, allows the compensation rationale to be extended from the individual decision to the social decision. The Kaldor-Hicks criterion is:

*"... a change from the present social state should be undertaken if the gainers from the change could compensate the losers in such a way that everyone would be better off, in that way satisfying the Pareto criterion."*⁵²

It is easy to criticize the Kaldor-Hicks criterion on moral grounds. However, the criterion holds up quite well when changes in the "social state" are small; the losers in one

⁵² Stokey & Zeckhauser, 1978, p.279.

change may be the winners in another. Thus, the Kaldor-Hicks criterion does provide a minimal rationale for the social decision-maker, namely that the polluting industries be required to invest in protection of those exposed to the carcinogenic risks, at the margin, up to an amount that would be sufficient to compensate the individuals in the exposed groups.⁵³

If it is assumed that all individuals in all the exposed groups have the same utility function and the same lifetime consumption, then the value of life for individual i is:

$$V_i = V(p_i)$$

The task is to allocate the regulatory budget by choosing among a set of differentiated standards; that is, different standards may be set for different exposed groups. For each standard j for each exposed group k , a mean value of life is required.

Let the distribution of risk under standard j across group k be $\phi_{jk}(p_i)$. Then the mean value of life for the group is:

$$V_{jk} = \int \phi_{jk}(p_i) * V(p_i) dp$$

Carcinogenic risks are usually modelled as lognormal distributions with many people exposed to small risks, and few exposed to high risks. Because $V(p_i)$ is convex in p_i , Jensen's inequality holds.

The mean value of life for the group > The value of life for the person with the mean risk.

⁵³ An alternative criterion could probably be developed using a Rawlsian social justice argument (Rawls, 1971). A member of each exposed group could be placed behind a veil of ignorance, and vote on the allocation of the regulatory budget. Each voter would know the risk distributions of all the groups, but not know to which group he or she belongs. The voting could be expected to be unanimous, if individual utility functions are the same. However, a Rawlsian criterion would be much more complicated to apply than the Kaldor-Hicks criterion.

The extension of the model to the social perspective is now complete. For any group that contains individuals bearing risks that are not “small”, then the average value of life for the group is greater than the small-risk value. To find out how much greater, we turn back to a consideration of the characteristics of the function $V(p)$.

The function $V = V(p)$ has two notable features -- it is flat for low risks, and it leaps suddenly to infinity when the risk is so high that the bearer could never be compensated. To apply the results, a small-risk value of life, V_s , and a limiting level of risk, p_{\max} , are required. Although both can be calculated, as shown in Figures 6 and 7, the calculated values are sensitive to lifetime consumption, degree of risk aversion, and, we may presume, the choice of utility function.

A small-risk value of life can be chosen by the social decision-maker from the empirical value of life literature; currently, the range is two to ten million dollars. For illustrative purposes, suppose that the small-risk value of life, V_s , is chosen to be five million dollars.

The limiting level of individual risk could be thought of as a *de intolerandis* bright line, that is, a risk that is in the region of intolerable risks.⁵⁴ The *de intolerandis* p_{\max} can be found through an exercise of introspection by social decision-makers themselves, or elicited from a randomly chosen sample of “ordinary people”.

⁵⁴ A risk that is *de intolerandis* is unthinkable; the Latin word carries a strong sense of obligation and necessity. The term *de intolerabilis* is not appropriate, since *intolerabilis* describes something that cannot be physically tolerated, such as a heavy weight. The word *incompensabilis* does not seem to exist, so a *de uncompensabilis* bright line is not an option.

The introspection process would involve mulling over the question:

What black pill risk is so large, that no amount of money would ever induce me to accept it?

The techniques developed for contingent valuation surveys could be used to help elicit answers. For instance, one could ask whether 20% was too high, or 5% too low. We might expect that the “average person” could not be induced to play Russian roulette for any amount of money; a risk of one in six, 17%, is too high. For illustrative purposes, suppose that the *de intolerandis* p_{\max} is chosen to be 10%.⁵⁵

The new (efficient and fair) decision heuristic for carcinogen regulation is now partly defined.

- Value the lives of individuals in the small-risk region at five million dollars.
- Value the lives of individuals bearing risks greater than 10% at infinity.
- For a group of individuals, use the mean value of life for the group. If the group contains any individuals bearing risks greater than 10%, then the mean value of life for the group is infinity. For such groups, the cost of protection is truly irrelevant; the risk is not to be tolerated. Bearers of risks that are *de intolerandis* must have these risks reduced to a tolerable level, irrespective of the cost.⁵⁶

⁵⁵ The choice of a *de intolerandis* black pill may be cognitively challenging, but the question is easier to answer than the questions that require answering for a *calculation* of the *de intolerandis* level. In order to calculate p_{\max} , an “average” utility function must be defined, requiring answers to a series of “standard gamble” questions.

⁵⁶ Note that the risk level of *all* the individuals in the group need not be reduced; engineering controls that shift the entire risk distribution to the left are not necessarily needed. Typically, the thin tail at the high end of the lognormal distribution would need to be truncated. If, for example, the carcinogenic risk is borne by the workers in a particular industry, the relatively few, who bear intolerable risks, could be paid a large wage premium to wear respirators.

Two missing components of the decision heuristic remain; both are concerned with the middle “in-between” risks, those for which the value of life will lie between the small-risk value and infinity. What is the upper bound of the small-risk region?⁵⁷ How does the value of life, $V(p)$, rise from five million dollars to infinity?

Turning back to the illustrative base case in Table 4 and Figure 5, we find a surprising result that yields answers to these two questions. The value of life $V(p)$ rises from a V_s of 4.87 million dollars extremely slowly, and then as the risk approaches p_{\max} , $V(p)$ suddenly rises very steeply, just before leaping to infinity. Over most of the “in-between” risk region, $V(p)$ rises exponentially, but at a very slow rate, and the leap to infinity is very abrupt.

In Figure 8, the value of life $V(p)$ is plotted against p . Figure 8 differs from Figure 5 in two ways. First, the value of life is plotted on the vertical axis, *not* the logarithm of the value of life. Second, the risk level stops at 3.5% in Figure 8, just short of p_{\max} , which is 3.57%.

It is now possible to see just how low $V(p)$ remains before its sudden leap to infinity. At $p = 3.5\%$, $V(p)$ is only twenty million dollars; at $p = 3.57\%$, $V(p)$ is infinite. Moreover, $V(p)$ does not rise above ten million dollars, until p exceeds 2.9 %. *Since the range for the value of life in the empirical literature is two to ten million dollars, it appears that virtually all risks are in the small-risk region, until the risk becomes intolerable.*

⁵⁷ A process of introspection could be used to find the upper bound of the small-risk region, but the results could not be taken seriously, since small probabilities are difficult to understand.




Figure 8. The solution to the Black Pill problem for the illustrative base case. The value of life $V(p)$ is shown varying with the size of the black pill, that is, the level of risk p . In the figure, the highest value of p is 3.5%, just short of p_{\max} , which equals 3.57%. Note that $V(p)$ ranges between five million dollars and twenty million dollars.

The new (efficient and fair) decision heuristic for carcinogen regulation is now complete, and can be expressed as follows.

- Risks are either small or intolerable.
- If risks are small, the benefits of reductions should be monetized, using a value of life taken from the empirical literature on preferences for lifesaving, and balanced against the costs of the risk reductions.
- If risks are intolerable, they must be reduced to a tolerable level, regardless of cost.

This result has led to the proposal of a new bright line of individual risk -- the concept of a *de intolerandis* level of risk. Unlike the *de minimis* and *de manifestis* bright lines, it is theoretically based, and can be measured through empirical studies. But it is a surprising and extreme result. Perhaps it is a vagary of the choice of utility function. In the next section, this possibility is examined.

Searching for a Less Surprising Result

The analysis in this paper has yielded a surprising result. The response of the value of life to the level of individual risk in the model is unexpectedly extreme; the size of the black pill is almost irrelevant, until it suddenly becomes all-important. The next task is to assess the strength of this result. A simple exponential utility function has been used in the model, and a different utility function may lead to a less extreme result.

In this section, a search is made for other utility functions that could be used in the Black Pill model. First, the properties of such utility functions are identified. Second, a class of utility functions that appear to possess these properties, and represent a more realistic attitude to risk, is examined.

Five Restrictions on the Utility Function

In the Black Pill problem, the individual is indifferent between rejecting the black pill of size p , or accepting the black pill for a compensation payment x .⁵⁸

$$U(W) = (1 - p) * U(W+x) + p * U(0) \quad (2)$$

$$p = [U(W+x) - U(W)] / [U(W+x) - U(0)] \quad (3)$$

The *first* required property of $U(W)$ is that, as the compensation payment (x) becomes infinite, p must tend to a limit that is less than unity. Otherwise, there must be some amount of money that is sufficiently large, which will compensate the individual for certain death. In other words, it would be possible to bribe people to kill themselves, which is not true (one hopes) for the average person.

As $x \rightarrow \infty$, $p \rightarrow p_{\max}$, and $p_{\max} < 1$, and equation (3) becomes:

$$p_{\max} = \lim_{x \rightarrow \infty} [U(W+x) - U(W)] / [U(W+x) - U(0)] < 1$$

$$p_{\max} = [U(\infty) - U(W)] / [U(\infty) - U(0)] < 1 \quad (5)$$

If $U(W)$ is not bounded from above, then $U(\infty) \rightarrow \infty$, and:

$$[U(\infty) - U(W)] / [U(\infty) - U(0)] \rightarrow \infty / \infty = 1$$

⁵⁸ The two simplifying assumptions -- zero consumption is tantamount to being dead, and no utility from bequests -- are not relaxed in what follows.

Thus, for p_{\max} to be less than 100%, $U(W)$ must be bounded from above.

This is a potentially important serendipitous discovery for general utility theory. Utility functions that are not bounded from above imply that it is possible to bribe people to kill themselves. Thus, any analysis using a von Neumann-Morgenstern utility function without an upper bound, and whose results rely on wealth tending toward infinity, is suspect. The number of analyses that fall into this category is a question that cannot be addressed here, but note that the analytically tractable (and, therefore, popular) logarithmic utility function is not bounded from above.

The *second* required property of $U(W)$ is that it must lead to a positive small-risk value of life, V_s , otherwise the lives of some people are being valued at zero. This property can be specified by finding an analytic expression for V_s .

By definition, $V = x / p$. As p tends to zero, x also tends to zero, so V tends to $0/0$, which is indeterminate.

Differentiate equation (2) with respect to x :

$$U(W) = (1 - p) * U(W+x) + p * U(0) \quad (2)$$

$$0 = U'(W+x) - p * U'(W+x) - [dp/dx] * U(W+x) + [dp/dx] * U(0) \quad (6)$$

Rearrangement of equation (6) gives an expression for dp/dx , and using L'Hopital's Rule:

$$\text{As } x \rightarrow 0 \text{ and } p \rightarrow 0, \quad V_s = \lim x / p = \lim [dx/dx] / [dp/dx] = 1/[dp/dx]$$

$$V_s = [U(0) - U(W)] / [-U'(W)] \quad (7)$$

Then $V_s > 0$ implies that $U(0)$ must be finite, placing a second restriction on $U(W)$.⁵⁹

Two properties of utility functions that will lead to solutions of the Black Pill problem have now been identified.

- If p_{\max} is to be less than unity, then $U(W)$ must be bounded from above.
- If V_s is to be positive, then $U(0)$ must be finite.

Any ordinal utility function that has these properties can, with the addition of scaling factors, be constrained to lie between -1 and 0 , so that:

$$U(0) = -1$$

$$U(\infty) = 0$$

The analytic expressions derived for p_{\max} and V_s -- equations (5) and (7) -- can now be simplified.

$$p_{\max} = -U(W) \quad (5')$$

$$V_s = [1 + U(W)] / U'(W) \quad (7')$$

Rearrangement of equations (5') and (7') gives the following analytic expressions for $U(W)$ and $U'(W)$.

$$U(W) = -p_{\max} \quad (8)$$

$$U'(W) = [1 - p_{\max}] / V_s \quad (9)$$

The utility at the initial consumption W , and its rate of change at the initial consumption W ,

⁵⁹ The simple logarithmic utility function, $U = \ln W$, does not possess this property either, since $U(W)$ tends to $-\infty$, when W tends to zero. However, $U = \ln(W+\alpha)$ is finite when $U = 0$, provided $\alpha > 0$.

are exactly determined by p_{\max} and V_s . Since p_{\max} and V_s are set exogeneously, equations (8) and (9) place a *third*, and a *fourth*, restriction on the utility function.

Finally, a *fifth* requirement for the utility function is imposed that would make it superior to the utility function, $U(W) = -\exp\{-W/\rho\}$ used in the illustrative solution to the Black Pill model. The simple exponential utility function has the property of constant risk aversion. Risk aversion is measured with a coefficient:

$$r(W) = -U''(W) / U'(W) \quad (10)$$

and, if $U(W) = -\exp\{-W/\rho\}$, $r(W) = 1/\rho$ for all W .

However, constant risk aversion over consumption is an unlikely attitude to risk.

*"It seems to be empirically true for many people that as their assets increase, they are only willing to pay a smaller risk premium for a given risk. Their reasoning is that as they become richer, they can better afford to take a specific risk, and therefore will forgo less to avoid it. The same reasoning implies that the insurance premium associated with an unfavorable lottery (i.e. one less preferable than the status quo) decreases as we get richer and increases as we get poorer."*⁶⁰

Hence, decreasing risk aversion may be a more natural property to impose on the utility function. Furthermore, the Black Pill model may yield a different solution, if it is solved using a utility function that has the property of *decreasing risk aversion*.

⁶⁰ Keeney & Raiffa, 1976, p.166.

A Candidate Class of Utility Functions

Five restrictions have now been set on alternatives to the simple exponential utility function. The first two are required properties, the third sets the value of the function at initial consumption, the fourth sets the slope of the function at initial consumption, and the fifth would allow for a more realistic attitude toward risk. The five restrictions are:

- The utility function must be bounded from above. Without loss of generality, the upper bound is set at zero.
- Utility must be finite at zero consumption. Without loss of generality, $U(0)$ is set at -1 .
- The value of the utility function at the initial consumption level is defined exogeneously by the choice of p_{\max} .
- The slope of the utility function at the initial consumption level is defined exogeneously by the choice of p_{\max} and V_s .
- The utility function must be decreasingly risk-averse.

A list of common decreasingly risk-averse utility functions contains only one that is bounded from above, namely:

$$U(W) = -\exp\{-a * W\} - b * \exp\{-c * W\} \quad \text{where } a, b, c > 0. \quad ^{61}$$

⁶¹ Keeney & Raiffa, 1976, pp.169-173.

In order to meet the two normalization conditions, a fourth parameter must be introduced. The utility function:

$$U(W) = -\alpha * \exp\{-\beta * W\} - \gamma * \exp\{-\delta * W\} \quad (11)$$

will be decreasingly risk-averse, with $U(0) = -1$, and $U(\infty) = 0$, provided $\alpha, \beta, \gamma, \delta > 0$, and $\alpha + \gamma = 1$.

More generally, utility functions of the form:

$$U(W) = -\sum \alpha_i * \exp\{-\beta_i * W\}, \text{ with all } \alpha_i, \beta_i > 0, \text{ and } \sum \alpha_i = 1$$

are bounded above, and have coefficients of risk aversion that decrease as W increases.

A utility function of this class may yield a solution to the Black Pill problem that is less extreme than the solution obtained by the use of a simple exponential utility function; the value of life function, $V = V(p)$, may rise faster with increasing p , and approach infinity as p nears p_{\max} with less alacrity. This possibility can be explored by examining the simplest function of this class, the weighted sum of two exponential functions in equation (11), since the risk aversion of a weighted sum of any number of exponential functions can be approximated by a function composed of the weighted sum of just two exponential functions.

The task is then to search for values for the four parameters α, β, γ , and δ , which lead to $V(p)$ behaving in a less extreme manner. If such values can be found, then there will be an argument for defining an “in-between” region of risks, over which $V(p)$ rises from V_s to infinity. For risks in the “in-between” region, the value of life should vary with the size of the risk.

If the rate at which $V(p)$ rises is maximized when p is small, then this will enable the largest possible “in-between” risk region to be identified. Examining the partial derivative of $V(p)$ with respect to p , for small p , shows that $V(p)$ rises most rapidly when the coefficient of risk aversion is large. The coefficient of risk aversion, $r(W)$ is equal to $-U''(W) / U'(W)$, (equation (10)), and $U'(W)$ is constrained to be constant (equation (9)). Thus, maximizing $r(W)$ is equivalent to maximizing $U''(W)$. This makes intuitive sense. In the illustrative case of the Black Pill model, the value of life begins to rise as $U(W+x)$ approaches the upper bound. In other words, the value of life rises when the slope of $U(W+x)$, that is, $U'(W+x)$, begins to fall rapidly.

Thus, the task of finding feasible values for α , β , γ , and δ , that will lead to $V(p)$ rising as fast as possible, can be set up as a constrained optimization problem. The objective function is $U''(W)$. Initial consumption, W , the small-risk value of life, V_s , and the *de intolerandis* level of risk, p_{\max} , are set exogeneously. Beside the requirement that α , β , γ , and δ all be positive (for decreasing risk aversion), there are three constraints. *

Suppose that W is set at 600,000 dollars, that is, 0.6 million dollars, V_s is set at five million dollars, and p_{\max} is set at 10%. Then equations (8) and (9) give:

$$-\alpha * e^{-\beta * 0.6} - \gamma * e^{-\delta * 0.6} = -0.1 \quad (12)$$

$$\alpha * \beta * e^{-\beta * 0.6} + \gamma * \delta * e^{-\delta * 0.6} = [1 - 0.1] / 5 \quad (13)$$

A third constraint on the four parameters, α , β , γ , and δ , comes from normalizing $U(W)$, so that $U(0) = -1$.

$$\alpha + \gamma = 1 \quad (14)$$

Optimal solutions for this problem lead to either β or δ being made vanishingly small, and the utility function in equation (11) loses the property of decreasing risk aversion. To avoid this result, a fourth constraint is added.

$$\delta \geq \epsilon, \text{ and } \epsilon \text{ is set} = 0.000001 \quad (15)$$

Values for α , β , γ , and δ are now found, so that equation (11) becomes:

$$U(W) = -0.932 * \exp \{ -5.617 * W \} - 0.068 * \exp \{ -0.000001 * W \} \quad (16)$$

This utility function has all the required properties, including a coefficient of risk aversion that decreases as consumption increases.

The results of using the utility function given by equation (16) in the Black Pill model are shown in Table 5. The value of life $V(p)$ rises to six million dollars when p is 1%, to ten million dollars when p is 2.7%, and above twenty million dollars when p is 3.36%. Thus, the “in-between” risk region lies between about 3% and the p_{\max} of 10%.

These values for α , β , γ , and δ are feasible, but are they plausible? An examination of the variation in the compensation payment, x , with the size of the Black Pill risk, p , suggests that the utility function in equation (16) is unrealistic. In Table 5, it can be seen that a payment of two million dollars and a payment of 100 million dollars are both acceptable as compensation for bearing a risk of 3.44%. Note the two decimal places in the risk of 3.44%. Compensation of one million dollars is acceptable for a risk of 3.43%, and compensation of a billion dollars is acceptable for a risk of 3.45%. The lower bound of the “in-between” risk region of 3% appears to be the result of a very unusual set of preferences.

Table 5. Testing a utility function with decreasing risk aversion.

The implicit value of life from the simple black pill model.

Parameters have been chosen so that the implicit value of life rises very fast.

The utility function is given in equation (16).

WTA compensation payments x (\$)	Black pill % probability of death p	log(V)	Implicit value of life (million \$) V
1	0.00002	6.70	5.0
10	0.0002	6.70	5.0
100	0.002	6.70	5.0
10,000	0.19	6.71	5.2
50,000	0.86	6.76	5.8
65,000	1.08	6.78	6.0
100,000	1.51	6.82	6.6
200,000	2.35	6.93	8.5
270,000	2.70	7.00	10.0
400,000	3.09	7.11	13.0
600,000	3.32	7.26	18.1
675,000	3.36	7.30	20.1
1,000,000	3.43	7.47	29.2
2,000,000	3.44	7.76	58.2
10,000,000	3.44	8.46	291
100,000,000	3.44	9.46	2,908
1,000,000,000	3.45	10.46	29,026
10,000,000,000	3.51	11.45	285,049
100,000,000,000	4.10	12.39	2,436,931
1,000,000,000,000	7.69	13.11	12,999,857
10,000,000,000,000	10.00	14.00	100,002,736
100,000,000,000,000	10.00	15.00	999,999,589

Since, the utility function in equation (16) is derived from requiring that the value of life rise as fast as possible, it appears to be impossible to find a utility function with the required properties, which will yield a lower bound for the “in-between” risk region that is less than about 3%. Indeed, the lower bound of the “in-between” risk region would approach the *de intolerandis* level as the utility function becomes more plausible.⁶²

Further experimentation with the class of utility functions with the desired properties may lead to the identification of plausible values for α , β , γ , and δ , and to a lower bound for an “in-between” risk region that is somewhat less than the *de intolerandis* level. In practice, however, there are likely to be few people exposed to risks that are above the small-risk region. A pragmatic modification of the policy implication would be to simply define the *de intolerandis* bright line a little lower than the value found empirically.

⁶² Note that the lower bound of the “in-between” risk region is the upper bound of the small-risk region.

Conclusion

Carcinogen regulation in the United States is characterized by conflict and strategic behavior. On one side, the Office of Management & Budget wants to set differentiated standards, in which levels of protection would vary across groups, and regulatory choices would be determined by the costs of saving a life. On the other side, the regulatory agencies hold to the principle of acceptable levels of individual risk, known as bright lines. Because explicit recognition of the conflict may allow for the creation and exploration of solutions that are inconceivable in the current entrenched positions, this paper begins with the placing of this conflict into Raiffa's negotiation framework. The negotiation analysis leads to a focus on the decision heuristics of the two parties. It is proposed that a new decision heuristic, in which averted cases of cancer are weighted by the distribution of the risk burden, should satisfy the normative concerns of both the OMB and the regulatory agencies. The specification of this heuristic leads to some insights.

The underlying rationale used for the new (efficient and fair) decision heuristic for the regulation of carcinogens is the principle of the payment of fair monetary compensation to the bearer of an involuntary risk. The Kaldor-Hicks criterion, that compensation need be only potential, allows the extension of this rationale to social decisions concerning risk. The compensation principle can be made operational by allowing the implicit value of life to vary with the magnitude of the risk borne.

Using a simplified variant of the Black Pill model developed by Howard, the implicit value of life is traced out as a function of the risk, and has several features that seem intuitively sensible and fair. The value of life is constant over a large small-risk range, but

eventually begins to rise, and finally, at some probability of death less than unity, leaps to infinity. This limiting probability represents a risk so large, that no amount of monetary compensation would make it tolerable; it is a *de intolerandis* bright line.

A disconcerting feature of the result is the extraordinary abruptness, with which the value of life leaps to infinity from the constant small-risk level. The risk must rise almost to the *de intolerandis* level before the value of life rises significantly above the small-risk value.

It may be that this feature of the result will change, if some of the assumptions used in the Black Pill model are relaxed. In particular, the use of a simple exponential function (constant risk aversion) representing the utility of consumption is questionable, since it is likely that risk aversion decreases as consumption increases. The properties of utility functions required for the generation of a value-of-life function with the desirable features of a constant small-risk value of life, and a *de intolerandis* probability less than 100% (certain death), are identified.

One of these properties is that the utility function must have an upper bound. Only a utility function that is bounded from above will lead to a *de intolerandis* probability less than unity. This implies that an individual with a utility function that lacks an upper bound can be compensated with some finite amount of money for committing suicide. The analytically tractable (and, therefore, popular) logarithmic utility function is not bounded from above.

A class of utility functions, which have the required properties, including decreasing risk aversion, is considered as an alternative to the simple exponential function. Functions in this class can be constructed that will lead to less extreme variation in the implicit value of

life. This leads to the identification of an “in-between” risk region, in which risks are neither small nor intolerable; in this region, the implicit value of life lies between the small-risk value of life and infinity. However, it seems unlikely that a lower bound for this “in-between” risk region, which is significantly different from the *de intolerandis* probability, exists, if consumption preferences are plausible. Even if such a lower bound exists, the development of a weighting function for the value of life, for the few people likely to be exposed to risks in this region, is almost certainly an unjustified policy refinement.

The implication of the result for public policy is the following. The regulatory agencies should identify individuals who are exposed to total involuntary carcinogenic burdens, that are above a *de intolerandis* bright line. Such risk burdens should be reduced, regardless of cost. For all others exposed to carcinogens of industrial origin, the benefits of risk reduction should be monetized using a constant small-risk value of life, and balanced against the costs.

This policy prescription is almost a corner solution -- in the “corner” of the OMB -- and appears far from the balanced negotiated compromise envisaged at the outset of this analysis. In part, this is a consequence of setting a hypothetical regulatory budget as a joint constraint on both parties. Under a regulatory budget, the normative prescription of the efficient allocators formalizes readily, since this is well-trodden economic territory. However, when the normative prescription of the bright liners is forcibly formalized into a ranking algorithm, its shortcomings become evident.

The placing of the conflict between the OMB and the regulatory agencies into a Raiffa-type negotiation framework, takes the disagreement to be over the objective of

carcinogen regulation. In reality, the conflict may have other causes. One may be that the two parties have different perceptions of the facts; another may be rivalry between the parties. It could be argued that the legal concepts of *de minimis* and *de manifestis* levels of risk reflect a concern with fair process, not with fair outcomes.

The negotiation in this analysis has served as a vehicle for exploring decision heuristics and objectives. In reality, OMB would be satisfied with the setting of regulatory budgets, and would leave the allocation to the regulatory agencies. The real negotiation would need to occur in the design and setting of the budgets.

However, in the absence of regulatory budgets, the main result of this analysis still holds. The concept of a *de manifestis* individual risk -- an acceptable level of risk from a single carcinogen -- has no underlying rationale, and should be abandoned. Instead, the concept of a *de intolerandis* individual risk -- an intolerable total carcinogen burden -- has an underlying rationale of fair compensation for risks involuntarily borne, and should be adopted.

More work is required to test this radical proposal. There may be some set of realistic decreasingly risk-averse utility functions, which will lead to a lower bound for an "in-between" risk region that is significantly different from the *de intolerandis* level. The addition of a bequest function, or the modification of the Howard model to allow for the latency period of carcinogenesis, might also change the policy implication.

The two key features of interest in the value of life function are the small-risk value of life and the *de intolerandis* level of risk. There is a large body of empirical work aimed at finding the former, and there is general agreement (among economists, at least) that the

small-risk value of life lies between about two and ten million dollars. A similar effort could be directed at eliciting the *de intolerandis* bright line. If the *de intolerandis* probability of death is found to lie in a narrow range, this would corroborate its value as a decision criterion.

References

- Breyer, S. 1982. Regulation and its Reform. Cambridge, Massachusetts: Harvard University Press.
- Broder, I. E. and J. F. Morrall. 1983. The Economic Basis for OSHA's and EPA's Generic Carcinogen Regulations. In What Role for Government? Lessons from Policy Research, ed. R. J. Zeckhauser and D. Leebaert:242-254. Durham, North Carolina: Duke University Press.
- DeMuth, C. 1984. OMB Position on Use of Risk Assessment, Cost-Effectiveness Analysis, Benefit- Cost Review in Setting Standards for Toxic Air Pollutants. Environment Reporter.
- Finkel, A. M. 1990. A Way Out of the "Individuals versus Populations" Dilemma in Air Toxics Regulation. Unpublished paper.
- Finkel, A. M. 1990. A Simple Formula for Calculating the "Mass Density" of a Lognormally Distributed Characteristic: Applications to Risk Analysis. Risk Analysis 10 (2): 291-301.
- Fisher, A., L. G. Chestnut, and D. M. Violette. 1989. The Value of Reducing Risks of Death: A Note on New Evidence. Journal of Policy Analysis and Management 8 (1): 88-100.
- Fisher, R., W. Ury, and B. Patton. 1991. Getting to Yes. New York: Penguin Books.
- Graham, J. D. 1995. A Historical Perspective on Risk Assessment in the Federal Government. Toxicology 102:29-52.
- Graham, J. D. and J. B. Wiener, eds. 1995. Risk versus Risk: Tradeoffs in Protecting Health and the Environment. Cambridge, Massachusetts: Harvard University Press.
- Haigh, J. A., D. Harrison, and A. L. Nichols. 1984. Benefit-Cost Analysis of Environmental Regulation: Case Studies of Hazardous Pollutants. Harvard Environmental Law Review 8: 395-434.
- Hammit, J. K. 1990. Probability is All We Have: Uncertainties, Delays, and Environmental Policy Making. New York and London: Garland Publishing, Inc.
- Howard, R. A. 1980. On Making Life and Death Decisions. In The Principles and Applications of Decision Analysis, ed. R. A. Howard and J. E. Matheson:483-506: Strategic Decisions Group.
- Howard, R. A. 1984. On Fates Comparable to Death. Management Science 30 (4): 407-422.

- Jones-Lee, M. W. 1989. The Economics of Safety and Physical Risk. Oxford: Blackwell.
- Keeney, R. L. and H. Raiffa. 1976. Decisions with Multiple Objectives: Preferences and Value Tradeoffs. New York: John Wiley & Sons.
- Linnerooth, J. 1979. The Value of Human Life: A Review of the Models. Economic Inquiry 17 (1): 52-74.
- Litan, Robert E. and William D. Nordhaus. 1983. Reforming Federal Regulation. New Haven: Yale University Press.
- Milvy, P. 1986. A General Guideline for Management of Risk from Carcinogens. Risk Analysis 6 (1): 69-79.
- Morrall, J. F. 1986. A Review of the Record. Regulation November/December: 25-34.
- Passell, P. 1994. Putting a Price Tag on Life, and Regulation That Protects It. The New York Times, November 10, C2.
- Portney, P. R. 1995. Cartoon Caricatures of Regulatory Reform. Resources 121: 21-24.
- Pratt, J. W. and R. J. Zeckhauser. 1996. Willingness to Pay and the Distribution of Risk and Wealth. Journal of Political Economy 4 (104): 747-763.
- Raiffa, H. 1982. The Art and Science of Negotiation. Cambridge, Massachusetts: Harvard University Press.
- Raiffa, H. 1996. Lectures on Negotiation Analysis: Kennedy School of Government, Harvard University.
- Rawls, J. 1971. A Theory of Justice. Cambridge, Massachusetts: Harvard University Press.
- Schelling, T. C. 1968. The Life You Save May Be Your Own. In Problems in Public Expenditure Analysis, ed. S. B. Chase: 127-176. Washington, D.C.: Brookings Institution.
- Sen, A. K. 1973. On Economic Inequality. Delhi: Oxford University Press.
- Stokey, E. and R. Zeckhauser. 1978. A Primer for Policy Analysis. New York: W. W. Norton & Co.
- Tengs, T. O., M. E. Adams, J. S. Pliskin, D. Gelb-Safran, J. E. Siegal, M. C. Weinstein, and J. D. Graham. 1995. Five-Hundred Life-Saving Interventions and Their Cost-Effectiveness. Risk Analysis 15 (3): 369-390.

Travis, C. C., S. A. Richter, E. A. C. Crouch, R. Wilson, and E. D. Klema. 1987. Cancer Risk Management: A Review of 132 Federal Regulatory Decisions. Environmental Science and Technology 21 (5): 415-420.

Travis, C. C., S. Richter Pack, and A. Fisher. 1988. Cost-Effectiveness as a Factor in Cancer Risk Management. Environment International 13: 469-474.

Travis, C. C. and H. A. Hattemer-Frey. 1988. Determining an Acceptable Level of Risk. Environmental Science and Technology 22 (8): 873-876.

Van-Houtven, G. and M. L. Cropper. 1996. When is a Life Too Costly to Save? The Evidence from U. S. Environmental Regulations. Journal of Environmental Economics and Management 30: 348-368.

Viscusi, W. K. 1993. The Value of Risks to Life and Health. Journal of Economic Literature 31: 1912-1946.

BELFER CENTER FOR SCIENCE AND INTERNATIONAL AFFAIRS

BCSIA is a vibrant and productive research community at Harvard's John F. Kennedy School of Government. Emphasizing the role of science and technology in the analysis of international affairs and in the shaping of foreign policy, it is the axis of work on international relations at Harvard University's John F. Kennedy School of Government. BCSIA has three fundamental issues: to anticipate emerging international problems, to identify practical solutions, and to galvanize policy-makers into action. These goals animate the work of all the Center's major programs.

The Center's Director is Graham Allison, former dean of the Kennedy School. Stephen Nicoloso is Director of Finance and Operations.

BCSIA's International Security Program (ISP) is the home of the Center's core concern with security issues. It is directed by Steven E. Miller, who is also Editor-in-Chief of the journal *International Security*.

The Strengthening Democratic Institutions (SDI) project works to catalyze international support for political and economic transformation in the former Soviet Union. SDI's Director is Graham Allison.

The Science, Technology, and Public Policy (STTP) program emphasizes public policy issues in which understanding of science, technology and systems of innovation is crucial. John Holdren, the STTP Director, is an expert in plasma physics, fusion energy technology, energy and resource options, global environmental problems, impacts of population growth, and international security and arms control.

The Environment and Natural Resources Program (ENRP) is the locus of interdisciplinary research on environmental policy issues. It is directed by Henry Lee, expert in energy and environment. Robert Stavins, expert in economics and environmental and resource policy issues, serves as ENRP's faculty chair.

The heart of the Center is its resident research staff: scholars and public policy practitioners, Kennedy School faculty members, and a multi-national and inter-disciplinary group of some two dozen pre-doctoral and post-doctoral research fellows. Their work is enriched by frequent seminars, workshops, conferences, speeches by international leaders and experts, and discussions with their colleagues from other Boston-area universities and research institutions and the Center's Harvard faculty affiliates. Alumni include many past and current government policy-makers. Libby Fellingner is BCSIA's Fellowship Coordinator.

The Center has an active publication program including the quarterly journal *International Security*, book and monograph series, and Discussion Papers for work in progress. Karen Motley is BCSIA's Senior Editor, and is also Deputy Editor of *International Security*. Members of the research staff also contribute frequently to other leading publications, advise the government, participate in special commissions, brief journalists, and share research results with both specialists and the public in a wide variety of ways.

BELFER CENTER FOR SCIENCE AND INTERNATIONAL AFFAIRS
JOHN F. KENNEDY SCHOOL OF GOVERNMENT
HARVARD UNIVERSITY

Discussion Paper Series

- 97-14 Anderson, Steven. "U.S. Electricity Restructuring and Environmental Policy: A Survey of Opinions in the Environmental Community."
- 97-13 Verma, Shashi. "Getting Privatization Right: Establishing a Competitive Market for Electricity Generation in Brazil."
- 97-12 Hahn, Robert. "Climate Change: Economics, Politics, and Policy."
- 97-11 Keohane, Nathaniel, Richard Revesz and Robert Stavins. "The Positive Political Economy of Instrument Choice in Environmental Policy."
- 97-10 Patt, Anthony. "Assessing Extreme Outcomes: The Strategic Treatment of Low Probability Impacts of Climate Change."
- 97-09 Long, Marybeth and Alastair Iles. "Assessing Climate Change Impacts: Co-Evolution of Knowledge, Communities and Methodologies."
- 97-08 Kandlikar, Milind and Ambuj Sagar. "Climate Science and Policy in India: Learning Some Lessons."
- 97-07 Franz, Wendy E. "The Development of an International Agenda for Climate Change."
- 97-06 Fisher-Vanden, Karen. "International Policy Instrument Prominence in the Climate Change Debate: A Case Study of the United States."
- 97-05 Agrawala, Shardul. "Explaining the Evolution of the IPCC Structure and Process."
- 97-04 Wright, Janice C. "'Bright Lines' and the Value of Life: Resolving the Dispute over the Regulation of Carcinogens."
- 97-03 Parson, Edward A. and Karen Fisher-Vanden. "Joint Implementation and its Alternatives."
- 97-02 Stavins, Robert. "Economic Incentives for Environmental Regulation."
- 97-01 Coglianese, Cary. "Assessing Consensus: The Promise and Performance of Negotiated Rulemaking."
- 96-03 Stavins, Robert. "Policy Instruments for Climate Change: How Can National Governments Address a Global Problem."
- 96-02 Stavins, Robert and Brad Whitehead. "The Next Generation of Market-Based Environmental Policy."
- 96-01 Mead, Andre and Sydney Rosen. "The Mixed Grass Prairie Reserve: Managing the Greater Badlands Region as a Whole System."

- 95-11 Victor Zaborsky. "Crimea and the Black Sea Fleet in Russian-Ukrainian Relations."
- 95-10 Robert Stavins, Adam Jaffee, Steve Peterson, and Paul Portney. "Environment Regulation and the Competitiveness of U.S. Manufacturing: What does the Evidence Tell Us?"
- 95-09 Olexity Haran. "Disintegration of the Soveit Union and the US Position on the Independence of Ukraine."
- 95-08 Graham T. Allison, Owen R. Cote, Jr., Richard A. Felkenrath and Steven E. Miller. "Avoiding Nuclear Anarchy: Containing the Threat of Loose Russian Nuclear Weapons and Fissile Material."
- 95-07 Richard A. Falkenrath. "The U.S.-Russian HEU Purchase Agreement: Achievements, Problems, Prospects."
- 95-06 Robert N. Stavins. "The Costs of Carbon Sequestrian: A Revealed Preference Approach."
- 95-05 William Dietrich. "The Challenge of Selecting Goals: Case Studies Regarding the Use of Critical Levels."
- 95-04 Michele N. Ferenz and Stephan D. Sylvan. "The Design of a Palestinian-Israeli Water Commision: A Best Practices Approach."
- 95-03 Robert N. Stavins and Tomasz Zylicz. "Environmental Policy in a Transition Economy: Designing Tradeable Permits for Poland."
- 95-02 Not available.
- 95-01 Brett A. Perlman. "Pricing the Internet: How to Pay the Toll for the Electronic SuperHighway."
- 94-10 Robert N. Stavins. "Correlated Environmental Uncertainty and Policy Instrument Choice."
- 94-09 Richard A. Falkenrath. "The United States and Ballistic Missile Defense after the Cold War."

The Belfer Center for Science and International Affairs (BCSIA) Discussion Papers, established in 1991, will be issued on an irregular basis with three programmatic subseries: International Security; Science, Technology, and Public Policy; and Environment and Natural Resources. Inquiries and orders may be directed to: Belfer Center for Science and International Affairs, Publications, Harvard University, 79 JFK Street, Cambridge, MA, 02138.

**BELFER CENTER FOR SCIENCE AND INTERNATIONAL AFFAIRS
JOHN F. KENNEDY SCHOOL OF GOVERNMENT
HARVARD UNIVERSITY**

Discussion Paper Series

- 94-08 Richard A. Falkenrath. "The United States, the Former Soviet Republics, and Nuclear Weapons: Problems and Policies of Denuclearization."
- 94-07 Robert Stavins and Adam Jaffee. "Environmental Regulation and Technology Diffusion: The Effects of Alternative Policy Instruments."
- 94-06 Not Available.
- 94-05 Victor Zaborsky. "Nuclear Disarmament and Nonproliferation: The Evolution of the Ukrainian Case."
- 94-04 Irving Lachow. "The Global Positioning System and Cruise Missile Proliferation: Assessing the Threat."
- 94-03 David Fairman and Michael Ross. "International Aid for the Environment: Lessons from Economic Development Assistance."
- 94-02 Robert W. Hahn and Carol A. May. "The Behavior of the Allowance Market: Theory and Evidence."
- 94-01 Rodd Bender, Wyman Briggs, Diane DeWitt. "Toward Statewide Unit Pricing in Massachusetts: Influencing the Policy Cycle."
- 93-06 Robert W. Hahn. "An Economic Analysis of Scrappage."
- 93-05 Bob Hancke. "Technological Change and Its Institutional Constraints."
- 93-04 Paul R. Portney and Robert N. Stavins. "Regulatory Review of Environmental Policy: The Potential Role of Health-Health Analysis"
- 93-03 Edward A. Parson and Richard J. Zeckhauser. "Equal measures and Fair Burdens: negotiating environmental treaties in an unequal world."
- 93-02 Robert N. Stavins. "Transaction Costs and the Performance of Markets for Pollution Control."
- 93-01 Rogier A.H.G. Holla. "Cultural Attitudes and Institutional Forces: Explaining National Differences in Biotechnology Policy."
-

The Belfer Center for Science and International Affairs (BCSIA) Discussion Papers, established in 1991, will be issued on an irregular basis with three programmatic subseries: International Security; Science, Technology, and Public Policy; and Environment and Natural Resources. Inquiries and orders may be directed to: Belfer Center for Science and International Affairs, Publications, Harvard University, 79 JFK Street, Cambridge, MA, 02138.

**BELFER CENTER FOR SCIENCE AND INTERNATIONAL AFFAIRS
JOHN F. KENNEDY SCHOOL OF GOVERNMENT
HARVARD UNIVERSITY**

Discussion Paper Series

- | | |
|-------|--|
| 92-13 | Charles H.W. Foster. "Of Vert and Vision: Ensuring the Legacy of the Northern Forest of New England and New York." |
| 92-12 | Gerald Jiro Hane. "Research and Development Consortia in Innovation in Japan: Case Studies in Superconductivity and Engineering Ceramics." |
| 92-11 | Raymond Vernon. "The Triad as Policymakers." |
| 92-10 | Dale W. Jorgenson and Peter J. Wilcoxon. "Energy, the Environment and Economic Growth." |
| 92-09 | Vicki Norberg-Bohm and William C. Clark, et al. "International Comparisons of Environmental Hazards: Development and evaluation of a method for linking environmental data with the strategic debate management priorities for risk management." |
| 92-08 | David Hart. "Strategies of Research Policy Advocacy: Anthropogenic Climatic Change Research, 1957-1974." |
| 92-07 | Matthias Dembinski. "Ballistic Missile Proliferation and the New World Order: A Critical Survey of the Literature." |
| 92-06 | Alexander I. Shlyakhter and Danial M. Kammen. "Estimating the Range of Uncertainty in Future Development from Trends in Physical Constraints and Predictions of Global Change." |
| 92-05 | Susan C. Eaton. "Union Leadership Development in the 1990s and Beyond: A Report with Recommendations." |
| 92-04 | Robert W. Hahn and Robert L. Axtell. "Reevaluating the Relationship between Transferable Property Rights and Command-and-Control Regulation." |
| 92-03 | Robert N. Stavins and Bradley W. Whitehead. "The Greening of America's Taxes: Pollution Charges and Environmental Protection," |
| 92-02 | Edward A. Parson. "Protecting the Ozone Layer: The Evolution and Impact of International Institutions." |

The Belfer Center for Science and International Affairs (BCSIA) Discussion Papers, established in 1991, will be issued on an irregular basis with three programmatic subseries: International Security; Science, Technology, and Public Policy; and Environment and Natural Resources. Inquiries and orders may be directed to: Belfer Center for Science and International Affairs, Publications, Harvard University, 79 JFK Street, Cambridge, MA, 02138.

**BELFER CENTER FOR SCIENCE AND INTERNATIONAL AFFAIRS
JOHN F. KENNEDY SCHOOL OF GOVERNMENT
HARVARD UNIVERSITY**

Discussion Paper Series

- 92-01 Lewis M. Branscomb. "S & T Information Policy in the Context of a Diffusion Oriented National Technology Policy."
- 91-15 Robert W. Hahn and Robert N. Stavins. "Economic Incentives for Environmental Protection: Integrating Theory and Practice."
- 91-14 Robert W. Hahn. "Government Markets and the Theory of the *N* th Best."
- 91-13 Edward A. Parson and William C. Clark. "Learning to Manage Global Environmental Change: A Review of Relevant Theory."
- 91-12 Lewis M. Branscomb. "America's Emerging Technology Policy."
- 91-11 Sean M. Lynn-Jones. "International Security Studies after the Cold War: An Agenda for the Future."
- 91-10 Rapporteur's Report of the Executive Session: "Negotiating a Global Climate Change Agreement." March 14-15, 1991.
- 91-09 Dale W. Jorgenson and Peter J. Wilcoxon. "Reducing US Carbon Dioxide Emissions: The Cost of Different Goals."
- 91-08 John D. Montgomery, et al. "Values in Conflict: Policy Interactions in the Pacific Basin."
- 91-07 Summary of Project 88/Round II Workshop Proceedings. "Incentive-Based Policies for Municipal Solid Waste Management." May 16, 1991.
- 91-06 Proceedings of Project 88/Round II Forum. "Market-Based Strategies for Environmental Protection: A Tribute to Senator John Heinz of Pennsylvania." May 16, 1991.
- 91-05 Charles H.W. Foster and Peter P. Rogers. "Rebuilding the Nation's Wetland Heritage: A Challenge for the 1990s."
- 91-04 Joshua Lerner. "The Impact of Patent Scope: An Empirical Examination of New Biotechnology Firms."

The Belfer Center for Science and International Affairs (BCSIA) Discussion Papers, established in 1991, will be issued on an irregular basis with three programmatic subseries: International Security; Science, Technology, and Public Policy; and Environment and Natural Resources. Inquiries and orders may be directed to: Belfer Center for Science and International Affairs, Publications, Harvard University, 79 JFK Street, Cambridge, MA, 02138.

**BELFER CENTER FOR SCIENCE AND INTERNATIONAL AFFAIRS
JOHN F. KENNEDY SCHOOL OF GOVERNMENT
HARVARD UNIVERSITY**

Discussion Paper Series

- 91-03 Darius Gaskins and Bruce Stram. "A Meta Plan: A Policy Response to Global Warming."
- 91-02 Gery E. Merchant. "Freezing CO₂ Emissions: An Offset Policy for Slowing Global Warming."
- 91-01 Adam Jaffee and Robert N. Stavins. "Evaluating the Relative Effectiveness of Economic Incentives and Direct Regulation for Environmental Protection: Impacts on the Diffusion of Technology."

The Belfer Center for Science and International Affairs (BCSIA) Discussion Papers, established in 1991, will be issued on an irregular basis with three programmatic subseries: International Security; Science, Technology, and Public Policy; and Environment and Natural Resources. Inquiries and orders may be directed to: Belfer Center for Science and International Affairs, Publications, Harvard University, 79 JFK Street, Cambridge, MA, 02138.