

Environmental Policy

What determines a person's political position on environmental issues? In chapter 6 we developed a theory of asymmetric "goodness" applicable to environmental issues as well as redistributive policy. A person is considered "good" if he supports environmental causes, but is not considered "good" if he opposes those causes. Group survival is the ultimate cause of that asymmetry. The long-term nature of the payoffs to environmental expenditures causes underinvestment in environmental amenities (from a group survival point of view) by a thoughtful democracy. In addition, the externalities of environmental amenities could produce goodness advocacy of more expenditures in an era without big government. With lags in determining good causes, environmental expenditures as a good cause could continue even with the externality corrections produced by big government.

In chapter 8 we found that those who had the greatest return from goodness were, indeed, those who supported environmental causes in addition to other causes with asymmetric goodness. In chapter 7 we saw that environmentalists engaged in more demonstrations than those opposed to environmental expenditures because the good demonstrate more than others. In this chapter we look for more evidence of asymmetric goodness for environmental issues. We also examine the policy consequences of that asymmetry in terms of positive economics.

The Phenomenon of Nonuse Value

There is strong evidence that some kinds of verbal behavior cannot be explained by the standard narrow self-interest model. Consider the literature on nonuse evaluation by environmental economics: where people are asked how much they are willing to pay (WTP) as their share of the costs to preserve some feature of the environment that they and their heirs will never use or see.¹ That literature is filled with controversy about whether such nonuse values are valid parts of the social

benefits of preserving environmental resources. But most agree that the answers cannot be explained by narrow self-interest. Those that believe in the importance of nonuse values often make their arguments in terms of altruism, or the inclusion of other entities' welfare in an individual's utility function.

The observed positive nonuse values to environmental amenities have an important property: asymmetry. There are both potential external benefits and costs when an individual successfully advocates for an amenity financed at public cost. The external benefits are others' use value of the amenity. The external costs are the costs or taxes that others incur because those who support a tax in favor of an amenity are supporting that tax for others as well as for themselves. If respondents to a questionnaire were simply using a cost-benefit assessment of the amenity and being altruistic, those external costs would be considered as well as the external benefits (Milgrom 1993). There is no evidence that users reduce their advocacy for the amenity in response to altruistic considerations toward nonusers. There is no reason a priori to expect this asymmetry in altruism.

Moreover, often the nonuse value assessed by nonusers is higher than the individual use value claimed by current and potential users. The required kind of altruism to fit such a picture gets extremely odd. To make nonuse value consistent with reasonable utility functions requires "planet love" or the inclusion of nonhuman welfare in the utility function. That goes beyond any altruism as normally defined to mean love for one's fellow humans rather than love for assorted environmental characteristics over and above the use of those amenities.² Those who believe that such an attitude exists would seem required to explain how it is consistent with evolutionary processes, since it would seemingly have nothing to do with either individual or group survival.

At first glance, it would appear at least conceivable that these estimated nonuse values could be produced by this expansive altruism that includes "planet love." However, the free-rider problem prevents either altruism or narrow self-interest from directly affect voting decisions. But the way nonuse values are estimated, a person is asked in effect, "If you were king, how much would you be WTP for an amenity if others also paid." His decision determines the hypothetical outcome. The free-rider problem appears to be avoided. Or has it?

The person knows that he is not king, that what he says in a survey will have even less impact on policy than his vote. Altruism cannot explain his survey answers as long as there is any private return from those answers.

There is, of course, a private return for claiming nonuse values: the desire to signal “goodness.” By asserting a WTP more for the amenity than its use value to them or even its value to potential users, people show that they are in favor of “good” causes, with the returns from that assertion previously discussed. There would be no similar payoff to concern about the taxpayers who bear the burden of environmental expenditures. The asymmetry of goodness explains the asymmetry of behavior between users and nonusers. That goodness is not free, however. It is constrained by the return to imitating the political positions of friends and one’s narrow self-interest.

What makes nonuse value so interesting is that there are so many ways in which it is inconsistent with utilitarianism—either narrow self-interest or altruism. Most of these ways have been summarized or developed by Diamond and Hausman (1993) and Diamond et al. (1993). We shall focus on some of their results. We add to their work in only two respects. The Diamond articles focused on nonuse value. But the behavior Diamond et al. found for nonuse value has far wider ramifications. They saw the connection to charity, but they did not explore the even more obvious connection to political behavior. What generates nonuse value generates a significant part of the demand for environmental legislation and for other “good” causes.

Second, Diamond et al. provided a convincing rejection of utilitarian explanations for nonuse values. But they did not provide a satisfactory alternative theory. Their alternative theory was “warm glow.” But, again, all warm glow means without further specification is non-altruism. Warm glow by itself does not predict that nonusers would get a warm glow by supporting environmental legislation, but that users would not get a warm glow opposing more expenditures for amenities. Warm glow must be more specific to yield such implications and other features of nonuse value. Our theory of asymmetric goodness does provide a sufficiently specified alternative to altruism to explain the behavior of nonuse value.

A consistent feature of nonuse values is that they increase little or not at all with increases in the size of the amenity in question. For example, as Diamond et al. showed, the amount people are WTP to save three specified wilderness areas is little more than the amount people are WTP to save any one of them. Different people are WTP roughly the same amount to protect two hundred thousand birds as two thousand birds. They also are WTP the same amount to prevent a decline in fishing in all Ontario lakes as to have the same effect on fishing in a subset of those lakes. These results are quite similar to the

findings of Palfrey and Prisbrey (1997) discussed in chapter 2. In their experiments net contributions to a public good do not increase with the productivity of the public good.

The embedding problem is a related finding from WTP studies. The amount people are WTP for an amenity is greater if they are asked separately how much they are WTP for that amenity than if that amenity is part of a list of amenities about which they are asked.³

Neither of these results makes sense as long as utilitarianism governs WTP. Our theory of charity, however, explains both results. The total amount of charitable contributions—where charity is broadly defined to include all prosocial acts—is determined by an individual's signaling needs and his conscience, an internalized form of signaling. He is roughly indifferent between charities that are equally satisfactory for signaling. Under those circumstances he makes little effort to discriminate between charities. In particular, he generally adopts the low-cost strategy of giving only to charities that seek him out. If, for example, he confined his total charity to protecting wilderness areas, he would give the same amount to protecting three wilderness areas as to one. In any case, a solicitation for a wilderness area must reduce the amount he is willing to give to any other wilderness area.⁴ From the point of view of charitable contributions, these two cases—greater size of the amenity to be protected and more causes from the same solicitor—are really the same case.

Desvousges et al. (1993) found another behavior inconsistent with WTP as a product of altruism. They found that when they asked about WTP in two different ways there was a consistent difference in the answers. They first asked people directly how much they were WTP to protect a given amount of waterfowl from oil spills (the open-ended form). They then asked others if they were WTP at least some amount for this protection, and then varied the amount (the dichotomous form). They found that these two procedures yielded very similar values of WTP at small and medium values of WTP. But there was a significantly greater percentage of people with high WTP values for the dichotomous form than for the open-ended form. This result is inconsistent with a utility-based WTP, which should produce the same WTP in both cases.

However, it is what we would expect if WTP is a signal. In chapter 3 we argued that a person's charitable contribution depends on that of associates. When others are choosing a reciprocity partner, they want the most trustworthy partner they can find. Hence, relative charitable contributions matter. In choosing their charitable contributions people

often want to know what is a reasonable amount of charity to give, that is, what others are likely to contribute. By asking whether one's WTP is greater than some large amount, the interviewer indicates to the respondent that that large amount is not a totally unreasonable amount. "Why, otherwise, would the interviewer bother to ask the question?" one queries. In contrast, the open-ended procedure provides the respondent no guide to a reasonable price. In consequence, we should observe, as we do, higher percentages of the WTP for the larger amounts given the dichotomous procedure.

"Above All Do No Harm"

Diamond and Hausman (1993) discuss a well-known paradox facing those who believe in a utilitarian explanation of contingent valuation.

Consider the issue of visibility at the Grand Canyon, recognizing how visibility varies throughout the year. Consider a costly project that can decrease pollution from power plants and thus improve visibility on some of those days. Next, consider a CV [contingency valuation] survey that asks respondents how much they are willing to pay (WTP) to fund this project to improve visibility. Instead of this survey, consider an alternative survey in which the respondents are told that the costly project has actually been approved (rather than just being proposed). Then tell the respondents that the government is considering saving money by canceling the project. In this alternative survey, the respondents are asked a *willingness-to-accept* (WTA) question: How much money would the respondents have to receive to be in favor of canceling the project (thereby accepting worse visibility)?

The two questions involve the same change in visibility. Thus one might reason that the two questions should receive the same answer, but, in fact, CV studies frequently find that WTA greatly exceeds WTP. (21)

Diamond and Hausman further show that this difference cannot reasonably be attributable to the most obvious explanation utilitarianism has to offer: the income effect. Goodness signaling, however, with a reasonable specification does provide an explanation.

Consider the design of mores to constrain individual self-interest in such a way as to maximize group survival. In some social interactions

a person benefits others. In other interactions a person harms them. Many of the beneficial social interactions can be accomplished with the minimum intervention of social rules. Trade or the reciprocal exchange of favors does the job. Harmful interactions are another story. Reciprocity will not work very well. In the absence of enforced social rules there are several strategies that a person can use to avoid harm. He can bribe somebody not to harm him. Unfortunately, this encourages threats of harm. Alternatively, a person can protect himself by counterthreats. But there will be many circumstances where it pays to make the threat a reality. Miscalculations can also occur, generating violence and counterviolence.

There is also a big difference in the side effects of social rules encouraging beneficence compared to social rules discouraging malevolence. Enforced beneficence produces the well-known disincentive effects of income redistribution. That mores are enforced by ostracizing rather than by the powers of the state should not change the direction of that effect. Proscriptions against harmful behavior reduce the resources required for either defensive or offensive behavior and, hence, tend to increase group survival. Therefore, we expect a far greater emphasis in mores against harmful behavior than in favor of beneficent behavior.

In consequence, it is a much more serious offense to violate the mores against harmful acts than to violate those in favor of beneficence. A person needs a greater compensating return to malevolence than she requires for not being beneficent.

WTP measures the worth of *increasing* an environmental amenity. WTA measures the worth of avoiding a *decrease* in the amenity. In the first case one is beneficent, in the latter case one is malevolent. Or is one? This characterization of the WTP and the WTA require asymmetric goodness about the environment. It is “good” to spend more for the environment; it is not “good” to save others the taxes required to finance the amenity. The difference between WTA and WTP is further evidence for asymmetric goodness on environmental issues.

Opaluch and Grigalunas (1991) and Boyce et al. (1992) argue that ethics generates the difference between WTA and WTP. However, they do not try to rationalize this moral value; they just state its existence. Kuran (1998) and Sunstein (1997) maintain that the WTP context forces the individual to focus on preferences and practical trade-offs, but WTA leads him to focus on the values he uses to evaluate preferences and choices. Socialized to consider it a moral obligation to preserve the environment in this latter case, the individual places less weight on his own preferences.

Though the details of their argument differ from ours, its logic has the same essential feature—*asymmetric goodness* in the social rules. The value system must place more emphasis on preserving the environment and less on supporting tax savings to others as well as oneself.

This discrepancy between WTA and WTP has wider ramifications. It would be much harder to rescind any proenvironment legislation than to prevent its enactment. It is harder to reduce a benefit to the poor than to prevent that benefit in the first place. Any effort to reduce tax rates to the rich are regarded as “redistribution to the rich” in spite of the progressiveness of the tax-benefit structure that would still exist even after such a reduction.

Of course, there is inertia associated with much legislation that increases the cost of change. This inertia generated by the goodness effect would, however, be an added source. One would expect more of it to be in evidence for issues involving *asymmetric goodness* than for other issues.

There is an alternative hypothesis that possibly could explain the difference between WTP and WTA—what Thaler (1980) called the *endowment effect* and Kahneman and Tversky (1984) called *loss aversion*. In a wide variety of experimental settings people’s utility for a state is increased when that state is ascribed to be the actual state. WTA is supposedly about how much one requires to give up an actual state, while WTP is about how much one is willing to pay to get the state, so this alternative hypothesis seems applicable to explaining the difference between the two.

Some of these results can be explained by nontrivial costs to switching consumption patterns. But some of the experiments focus on ownership rather than consumption. For example, people are reluctant to sell stock that they have inherited even though they are reluctant to buy the same stock with cash they have inherited and they are told that brokerage costs are trivial (Samuelson and Zeckhauser 1988).

To determine whether the *endowment effect* is really applicable to the difference in the WTA and WTP case it would be helpful to understand the reason for the *endowment effect*. Psychological decision costs might be the explanation. There can be a cost to our ego in making a wrong decision that is not fully compensated by ego returns from making a right decision. Look at the example cited in the last paragraph. In one case one decides whether to buy stock; in the other case one decides whether to sell the same stock. One can seemingly reduce the psychological decision costs in both cases by doing nothing. But this requires that the ego costs of the wrong decision to do nothing be

less than those costs of the wrong decision to do something. That condition would be satisfied if sins of omission are regarded less seriously than sins of commission even when there are no external consequences of those sins.

If that were the explanation for the endowment effect, it does not appear applicable to the difference in WTA and WTP cited by Diamond. There the government is either considering approving the environmental project or canceling the project. Neither of these decision processes is initiated by the respondents. They are confronted by decision costs in any case.

Environmental Federalism: Theory

Oates and Schwab (1988) looked at the regulation of environmental externalities confined to a locality. Given several reasonable simplifying assumptions, local governments will adopt efficient environmental standards. There will be no “race to the bottom” of environmental standards. The fiscal benefits from attracting capital by lowering standards below the efficient level will be more than offset by higher wages and reduced amenity levels.

There are important reasons why under these circumstances regulation should be localized, and they are related to Oates’s decentralization theorem (1972). They concern the greater ability of local regulation to respond to variation in local conditions as compared to federal regulation. The benefits from regulation and the preferences for these benefits are likely to vary by locality. The costs are also likely to vary. It would appear, then, that local regulation of “local externalities” is preferable to national regulation.

Nevertheless, there are many cases where regulation occurs for a wider area than the nature of the externalities justifies. This regulation also often imposes more stringent standards rather than less stringent standards than the local residents prefer. To our knowledge, nobody has asked why, much less provided an answer to the question. We review some pertinent evidence.

Such an answer is easy to generate given asymmetric goodness signaling for environmental issues. Suppose there are two geographic areas and all the benefits and costs of the “localized externality” are confined to one of these areas. If there were only localized regulation with voting by citizens of each separately, the residents of that area would choose a level of regulation consistent with those costs and benefits in addition to the goodness returns they get from voting for

proenvironmental causes. If there were regulation of both areas voted upon by the citizens jointly, the nature of the voting in the affected area would not change. But those in the unaffected area would not be indifferent. They get a return to signaling goodness with none of the costs associated with putting the regulation into effect. They will opt for a higher level of regulation than would be chosen by those in the affected area, even though they receive none of the benefits from the regulation. In the unaffected area the demand for environmental regulation is unlimited, since there are no cost constraints. In consequence, the two areas together will vote for more regulation than the affected area would prefer.

Furthermore, there is an incentive for the unaffected area to advocate regulation on the basis of the joint areas in order to get more opportunity to signal goodness. This might be sufficient to overcome the opposition of the affected area to nonlocalized regulation, especially if the population of the unaffected area is large relative to the population of the affected area. It is by no means certain, however, that two-area regulation will occur. However, if it does, it will impose stricter standards than localized regulation for "localized externalities."

There is another process that can lead to the same set of conclusions. It also depends upon goodness asymmetries. The analysis of Oates and Schwab assumed that local voters were motivated either by narrow self-interest or altruism confined to local borders. Goodness signaling changes their conclusions. The locality would vote for more environmental regulation than the utilitarian interests of its voters would dictate. This excessive amount of regulation will not generate a fully compensating reduction in wages because the value of the increased amenity is less than its costs. There will be some combination of capital and labor flight in response to a loss in real income produced by the excessive regulation. It is possible that those in the locality would show by their votes that they prefer national regulation of all localities to reduce this capital and labor flight. If in this case national regulation occurs, it will involve a higher level of regulation than would have been imposed locally. The prospective flight of resources would constrain local regulation in a way that it would not constrain national regulation.

There is one important difference between this second process and the one previously discussed. The second depends upon areas being similar; the first depends upon areas being dissimilar. Suppose that one locality is the only one that would be affected by a regulation. Then, making the regulation national would not stem the flight of capital or labor. It would simply make that flight greater by increasing the level

of regulation. There can be no local majority for nationalizing the regulation.

In contrast, if all areas were affected equally by the regulation, the first process would not work. There would be no areas where voters could costlessly signal their goodness because all areas would have to bear the costs of the regulation.

But whether the two areas were the same or different, we would get the same result—stricter regulation at the centralized level. Furthermore, that same result requires goodness asymmetry in both cases. When the areas are equally affected, more joint regulation is generated by the reduced flight of capital compared to localized regulation. This flight occurs only because the overregulation at even the one-area level implies that the reduction in wages does not fully compensate for the cost of regulation. When only one area is affected, the greater centralized regulation is produced by goodness advocates in the unaffected area.

The entire analysis of this section has made an assumption that is roughly appropriate for most environmental regulation. There is not an important redistributive component of the regulation. One does not need asymmetric goodness to explain the centralization of laws that are primarily redistributive in character.⁵ Centralization can be produced to avoid the movement of harmed people and capital out of a locality and the movement of the beneficiaries into the locality. This has been the usual explanation for national as opposed to local taxes. But even in this case “asymmetric” goodness contributes to this centralization.

Some cases where redistribution is involved are better understood in terms of asymmetric goodness than in the flight of resources—cases where resource flight is probably not very important. Take the demands of developed countries for restrictions on child or prison labor in less-developed countries. The products of this kind of labor in less-developed countries are usually not close substitutes for the products of developed countries. Therefore, the developed countries probably have more to gain in terms of lower prices from child labor than they lose from a flight of capital. Part of developed countries’ opposition to such labor can be attributed to the power of unions. But there are many nonunion opponents in developed countries to such labor in less-developed countries. Asymmetric goodness seems a required part of the explanation.

Similarly, the European Union requires of its member countries that they have no death penalty. Surely, there are no direct benefits or costs to people outside the country involved. Goodness must be operating, though in this case there could be an opposite morality signal.

Environmental Federalism: Evidence

Those in localities that are the primary beneficiaries of the benefits and bearers of the cost of regulation oppose the stricter environmental standards that others would impose. Kalt and Zupan (1984) analyze senatorial support on roll call votes for stricter standards associated with the Surface Mining Control and Reclamation Act (SMRCA) in 1977. One of their results: the higher the state's surface coal mining resources as a fraction of state income, the greater the opposition to stricter standards.

Durden, Shogren, and Silberman (1991) study votes in the House of Representatives in 1974 on support for controls of strip mining. They too find a significant negative effect of mining employment. Even without asymmetric goodness in their theoretical arsenal, neither set of authors was surprised by their results. Obviously, the cost of these stricter standards fall primarily on the localities in which surface mining is important. However, they fail to see that the benefits of these stricter standards also fall on these localities.

Nonlocals are benefited only to the extent that they visit the areas adjacent to the surface mines. For the purposes of either tourism or hiking, these surface-mining areas tend to have close substitutes. In consequence, nonlocals are unlikely to benefit much from the grooming of former coal mines. And nonlocals bear some of the cost of more expensive surface mining in the form of more expensive coal. We suspect these direct nonlocal costs are greater than the direct nonlocal benefits. Asymmetric goodness is required to explain these results.

Even without goodness it is conceivable that locals would favor national regulation. While a coal mine cannot move from area to area, the amount produced can shift. The reduced mobility of production with national regulation could make it desirable to the area affected. But that is not the case, as witness the local area's opposition to this regulation. Asymmetric goodness does operate.

Kahn and Matsusaka (1997) explain support for environmental initiatives in California during the period 1970–94. For a number of issues it is likely that all externalities were local. These included a 1982 vote on mandated bottle deposits, a 1990 vote on forest preservation and its counterinitiative, and a 1990 initiative to ban hunting of mountain lions. In nearly every case the coefficients of the variables meant to control for local residence were significantly negative in explaining support.

Mandated bottle deposits express a concern with unsightly trashing

of the countryside. Those visions are almost exclusively for local eyes. In the case of the forest preservation initiative the apparently nonlocal effects are not very important. Nonlocal hikers like to hike in forests, but nearly all that hiking from the outside occurs on public land, and the forest preservation proposals are only relevant to private land. There is a worldwide concern with the preservation of forests to reduce CO₂ levels in the air, but California forests would only have a trivial impact on that goal. The only people that are likely to see a mountain lion are locals. They, too, are the ones who pay the price for any mountain lion attacks on livestock, and are most likely to enjoy hunting the lions. Of course, some nonlocals would also like to hunt mountain lions. But in this case this nonlocal interest is at variance with the nonlocal goodness interest, and clearly cannot explain nonlocal opposition to hunting mountain lions.

Asymmetric goodness is the obvious explanation for these attempts to centralize decisions about these “localized externalities.” The evidence also indicates that the opposition to these initiatives was concentrated in the localities that would be affected by them—another prediction of asymmetric “goodness.” When the authors controlled for the percentage of employment in construction, the percentage of employment in farming or forestry had a significant negative effect on support for environmental preservation ten of twelve times and never had a significant positive effect. Again, people from outside the localities were attempting to impose stricter standards than the locals desired.

Dineen and Twail (1997) document another case of the federal government’s imposing minimum standards for a “localized externality.” Contamination of water systems by “adjusted gross alpha emitters” that are carcinogens likely to have entirely local costs because only long-term prolonged exposure puts people at risk and because the cleansing capacity of streams is sufficient to assure no downstream contamination if even minimal locally approved standards are enforced. Yet the federal government set minimum standards. Two hundred and eighty water systems failed to meet this minimum standard. For those localities, obviously, the locality preferred to do less than what the federal government required.

Dineen and Twail show that in the case of the water systems that failed to meet the standards, enforcement imposed substantial net costs on the localities, even if the benefits of a cancer prevented is assumed to be a very high \$10 million. Of course, this net cost is not

sufficient to insure that the localities would fail to meet the standard. Goodness operates on the local level as well as on the national level. It is just that its effects will not be as great locally. Their results, however, show that there is asymmetrical goodness for this environmental issue. Somebody must be getting some benefit from requiring higher standards than can be justified on utilitarian grounds.

In another case Morris (1997) examines federal regulations on pesticide use in agriculture. He documents that the Environmental Protection Agency restricts eradication programs against fire ants and predators that would have little external costs beyond state borders.

There is another case of “localized externalities”: animal trapping. The consequences of trapping animals such as beaver are focused almost exclusively in the localities in which they are being trapped. Where animals are not a tourist attraction, it is the locals who experience both the costs and benefits of having the animals around. A National Public Radio broadcast in 1999 indicated that the impetus for more stringent regulations on trapping comes from urban areas; the opposition comes from rural areas. These results make sense only in terms of asymmetric “goodness.”

Here is yet another case. The federal government sets aside more wilderness area in Alaska than Alaskans’ want, as evidenced by the behavior of their congressman and senators. For example, the Arctic National Wildlife Act was approved in 1979 by the House of Representatives by a vote of 360–65 and approved by the Senate in 1980 by 78–14 (*Congressional Quarterly* 1979, 1980). The entire Alaskan congressional delegation opposed this bill that set aside large wilderness areas in Alaska. Though there were only three Alaskan votes on this bill, that number is sufficient to reject the hypothesis that Alaskan votes were a random sample of all votes.⁶ Alaska’s isolation from the rest of the United States makes virtually all environmental regulation the regulation of “localized externalities.” The number of tourists going to Alaska from the rest of the United States is trivial compared to the U.S. population, and most of them are confined to a narrow maritime strip that is not affected by most wilderness area regulation. One expects the rest of the United States to be affected more by the effect of this regulation on Alaskan exports than on tourist opportunities. And the export price effect would discourage others from supporting more wilderness areas in Alaska. Asymmetric goodness seems required to explain the imposition of wilderness areas on unwilling Alaskans.

Cost-Benefit Analysis

As the name suggests, cost-benefit analysis simply sums up the costs and benefits of any policy to determine whether the policy is a good idea or not. In the “pure” form these costs and benefits are determined by private assessments as manifested through market behavior. The “pure” form also uses market interest rates to discount costs and benefits, where the market interest rate is defined as the interest rate facing investment alternatives (but, as we saw earlier, care must be taken to appropriately estimate the time stream of the benefits and costs).

An impure form of this analysis also includes nonuse value or uses interest rates lower than market rates (never higher). Only the “pure” form is consistent with utilitarianism. As we previously saw, nonuse values have no utilitarian meaning. Furthermore, the use of lower than market rates imposes time preferences other than what people want.

Cost-benefit analysis is utilitarianism at work. Most, but not all, of the criticism of cost-benefit analysis from environmentalists is a criticism of utilitarianism.⁷ They argue that environmental values are morally superior to consumer values and, hence, should not be evaluated simply by what consumers want. There is, of course, a difference between environmental values and consumer values in the sense that there are important externalities in the former. But the whole purpose of cost-benefit analysis in this context is to evaluate the externalities, not to ignore them. There has to be something else that gives environmental values their superiority.

Many of the critics go no further than this declaration of the virtues of environmental values, as if they were so obvious as to require no defense. Others provide assorted arguments, all of which are ultimately based on group survival or its natural misinterpretations that have been previously discussed. One of the arguments makes private consumption inferior to public consumption because of the selfish basis of the former (Sagoff 1988). This is a familiar refrain of “do-gooders” who ignore the virtues of selfishness when properly channeled.

Another argument focuses on the long-term benefits of environmental policy. Many environmental regulations have benefits over the generations with costs concentrated in the present. This argument would question the use of market interest rates as an appropriate discount factor for cost-benefit analysis. Indeed, as we have seen, maximizing group survival requires lower discount rates than man would freely choose.

Both of these arguments have one thing in common. They are non-utilitarian. They advocate people as a collective buying something that they do not want as individuals. The critics of cost-benefit analysis recognize what they are doing. They also argue that it should be these collective decisions rather than the private decisions of cost-benefit analysis that should count.

The critics are right in believing that these collective decisions will differ from cost-benefit decisions. They are also right in believing that the collective decisions will consistently favor greater environmental regulation than would be produced by cost-benefit analysis. The obvious explanation: it is "good" to be in favor of environmental expenditures by the government; it is "bad" to oppose them.

The beauty of the cost-benefit debate for our purposes is that it is another clear demonstration of a nonutilitarian component of environmentalism. It is also a demonstration of goodness asymmetry. Opponents of environmentalism are more than willing to use cost-benefit analysis in determining environmental policy. The 1994 Republican "Contract with America" proposed replacing public health mandates in with cost-benefit analysis and redoing all past regulation in light of cost-benefit analysis.

Such conclusions will hardly shock anybody except the economists who use a narrow self-interest model of political behavior, but that is, of course, the dominant view of economics. What is more interesting is our explanation for this phenomenon: asymmetric "goodness." It is hard to understand how else moral values would arise that have nothing to do with the costs or benefits to the individuals producing those values.

In discussing the critics of cost-benefit analysis we are not just dealing with radical extremists. This criticism, or the preferences that generated it, is in the mainstream, so much so that it has had a profound impact on public policy. Cost-benefit analysis, so obvious from a utilitarian perspective, has not won the day in determining environmental policy. "Congress has treated environmental risks as impermissible except when required by considerations of feasibility. Rather than cost-benefit analysis, Congress has adopted a proenvironmental baseline for the control of air and water pollution, carcinogens in the workplace, and hazardous waste sites, and has much less often called for cost-benefit analysis" (Farber 1999). For example, the Clean Air Act explicitly rejects considerations of cost in determining the appropriate level of air quality. And even where cost-benefit analysis is used, it has been distorted by President Clinton's executive order to allow "contin-

gent valuation” (nonuse value) in its calculation. As we saw earlier, nonuse values have no utilitarian base. The underlying problem is that goodness has so permeated popular opinion that cost-benefit analysis is not a winning cause. That is why the critics of cost-benefit analysis would prefer political to economic judgments about environmental policy.

The Value of Life

The utilitarian rationale of most environmental regulation is to protect health when the market fails to do so either because of externalities or lack of information of market participants. What regulations are desirable? The cost-benefit answer to this question is to compare the cost of the health benefits from government regulation to the cost of the same health benefits implicit in market behavior. Practitioners tend to focus on one component of those health benefits—the value of life. Obviously, such a focus has its problems. Health affects the quality of life as well as life expectancy. However, that is true for market-determined health as well as government decisions. We do not know of any reason for a systematic difference between the ratio of death to ill-health generated by the two classes of decisions. Hence, we suspect that the order of magnitude of that ratio is the same for the two decisions, and that is all that is required for our purposes.

However, there is one systematic difference between the impact of the government regulations we examine and market behavior. The Clean Air Act, for example, primarily prevented deaths from respiratory cancers, which tend to occur late in life. Many of the market-determined behaviors examined are related to deaths from injuries that occurred throughout life, but mostly in adolescence and young adulthood. We do not propose that one substitute a standard of years-of-life saved for the number of lives saved. The emotional and financial investments in very young children are substantially less than in older ones. But those investments have mostly already been made for young adults. Saving their lives must, certainly, be more valuable than saving lives of the average lung cancer patient. So the benefits of saving a life by market behavior probably exceed the value of saving a life under the Clean Air Act, possibly by a substantial amount.

Keeping these reservations in mind, we will follow the standard practice of cost-benefit analysis. It says that a regulation is better than no regulation if it saves a life at lower costs than does the market. If, then, the cost-benefit approach were the single principle governing

health-related environmental regulation, one would predict that absent mistakes, the value of life for all regulations would be lower than or equal to the market-generated value of life. Given mistakes in regulations, this proposition is not easy to test directly. However, it does imply that for any class of economic regulation, the expected value of life will be less than the expected market value of life.

For regulations under the Clean Air Act just the opposite occurred. Miller (1989) summarizes twenty-nine high-quality studies of the value of a life determined by market behavior. He finds in 1989 dollars that the mean value is \$2.25 million. Van Houtven and Cropper (1996) find that the mean value of a cancer prevented for fourteen banned uses under the Clean Air Act was \$348 million, also in 1989 dollars.

The concepts used in the two studies are not quite the same. A cancer prevented is not a life saved. Not all cancers are fatal, and the air pollution that produces cancers reduces life expectancy in other ways, such as increasing emphysema. These differences work in opposite directions. We estimate that at most the number of cancers saved should be increased by 40 percent to be equivalent to the number of lives saved under the Clean Air Act.⁸ That would mean that the average value of a life saved under the Clean Air Act would fall at the most to \$249 million, still substantially greater than the market-determined mean value of life. This large difference is statistically significant at the 1 percent level.⁹

This result is hardly surprising. The Clean Air Act prohibits the EPA from using cost-benefit analysis. However, the EPA seems to have somewhat violated that prohibition. Van Houtven and Cropper show that in twenty uses considered under the Clean Air Act but not banned and where banning could save lives, the mean value of a life saved was \$11,571 million, substantially more than the value of a life for banned uses. That result is also not surprising. Congressmen could signal their goodness by voting to prohibit cost-benefit analysis, but the enormous waste of resources involved in totally ignoring costs was too much even for the EPA.

There is an alternative explanation for the discrepancy between market-determined values of life and Clean Air Act values of life: lack of information by market participants. The relatively low standard deviation in market-determined values of life calculated in considerably different ways (\$.58 million in Miller's twenty-nine studies) suggests that this is not a terribly serious problem. The values of life under the Clean Air Act are significantly greater than the largest market-determined value of life. Furthermore, even if there were a serious

information problem, values of life saved by government regulations should be less than market-determined values of life if government regulations were determined simply by cost-benefit calculations. The regulations would focus on those areas where the information problem produced the greatest downward bias in market-determined values of life.

Goodness signaling has another implication. Since it is not about consequences, there is no reason for regulations in various areas to generate the same utilitarian consequences. In particular, the values of life should be quite different regulation to regulation. Of course, one would expect some of this just by the all-or-nothing character of government regulation. A use is either banned or not banned. It is sensible to ban higher value of life uses as well as lower values as long as both are less than the market value of life. But given the goodness motivation for banning, one would expect to see some uses that are not banned having lower values of life than uses that are banned. And such cases cannot be explained on utilitarian grounds. Some of this could be explained by mistakes. However, if one were to find a class of non-banned uses that have a significantly smaller value of life than a class of banned uses, the mistake hypothesis can be ruled out.

Van Houtven and Cropper (1996) find that the 149 unbanned uses considered by the EPA under the Federal Insecticide, Fungicide, and Rodenticide Act had a mean value of life of \$15.697 million, considerably less than the banned uses for the Clean Air Act—\$348 million. The published data do not permit a statistical test given the nonnormality of the distributions. But the large difference in means is suggestive. The most obvious explanation for the difference in results under the two different acts is that the Insecticide Act did not prohibit the use of cost-benefit analysis.

Animal Rights

In previous chapters we developed one big consequence of this nonintellectual approach to “goodness.” Often the goals of goodness will be understandably derived in an emotional sense from group survival, but they will not in fact contribute to that survival. From the point of view of group survival the culprit is misplaced compassion. In chapter 6 we saw this in the case for criminals, war victims, and women, among others. A similar compassion operates in the case of environmental policy: compassion toward animals. Such compassion is required to rationalize the Endangered Species Act, which cannot be defended either in

terms of cost-benefit analysis or maximizing the survival of humans as a group. Indeed, the Endangered Species Act specifically rejects consideration of costs except under very special circumstances.

The best man-oriented defense is Wilson's (1992). Diversity in DNA is potentially useful to man for medicines and other products, and there is user value in biological diversity. Having said this, however, he reveals his nonutilitarianism: "We should judge every scrap of biodiversity as priceless."

The DNA that has been found useful thus far comes to our knowledge exclusively from plants, not even remotely related to the animals that have been protected under the aegis of this act. The animals that man has found useful as models to test medicines are also not protected. There is evidence that the DNA argument is not the driving force behind this act. Given recent developments in biology, DNA can be preserved and multiplied without keeping the plant or animal alive. There has been no great movement to eliminate the Endangered Species Act on that account. And even without DNA preservation there are such things as zoos and botanical gardens that permit DNA preservation but that do not require large tracts of land to be set aside for that purpose.

It is quite likely that there exist species worth preserving in terms of costs and benefits appropriately calculated. But it is also likely that there are species that are not worth preserving by the same standard. A blanket protection for all species seems singularly inappropriate from a utilitarian perspective.

Environmentalists do not rest their case for the Endangered Species Act on specious utilitarian grounds alone. For example, Farber (1999), a moderate on environmental issues, assesses the general attitude of the population as a whole, including environmentalists.

Most people today recognize that nature has value, quite apart from any immediate utility. Even beyond aesthetic appeal, we can recognize that nature is the result of a process beyond human scale, whether in the form of divine intervention or the sheer extent of a billion years of evolution. Together with the more utilitarian reasons for preserving biodiversity to provide direct human benefits, these values deserve a place in our societal pantheon. (109–9)

There is at least a modicum of a utilitarian defense for the Endangered Species Act. It would be hard to explain other features of good-

ness behavior toward animals on pragmatic grounds. What utilitarian goal—when utilitarianism is confined to humans—would be achieved by the animal rights activists? Where is the gain that can compensate for the human losses that would be produced by restricting the use of animals for medical experiments?

Paraphrasing a frequent argument made by environmentalists and animal rights activists alike, “This universe consists of more than just humans. Other animals also have the right to live and thrive on this planet.” Indeed, there is no reason to suspect that this is a man-centered universe. But that is not a good argument for why man should not be man centered. We are talking about decision making by and for man, not by and for seals. From a survival or a utilitarian perspective there is no more reason for man to be concerned with seals than seals for man except for man’s joy in watching seals or wearing sealskin coats.

This widespread compassion toward animals is a recent phenomenon. It is associated with television’s making us aware of details in their lives, anthropomorphizing them, and claiming man’s “cruelty” in endangering their habitat. A substantial percentage of the nature programming on PBS and the Nature Channel has this as its theme. As discussed in the last chapter, when preferences were being developed, individual survival was enhanced by being compassionate toward friends because that compassion was reciprocated. Friends were people whose lives we know a lot about. That compassion has mistakenly been transferred to animals that we know something about. This process started with pets, who originally served utilitarian purposes, and has now extended to the animal kingdom in general.

The contents of this chapter have a special importance. We have developed a theory of asymmetric goodness based on group survival and its misinterpretations. There is a competitive theory, which supposes that big business dominates the political process (for example, Chomsky 1989) but does not dominate the development of mores. Then, we would also expect asymmetric “goodness.” (If big business dominated both, then political outcomes and goodness would be the same.) It would be “good” to be opposed to big business. But this kind of goodness would have one big difference from the goodness we have discussed up to now. The new goodness would be based on what people *want* that big business is preventing them from getting. In consequence, it would be utilitarian in nature. But in this chapter we have seen that the goodness ethic is the antithesis of the utilitarian ethic. This provides additional support for our theory of the origins of asymmetric “goodness.”