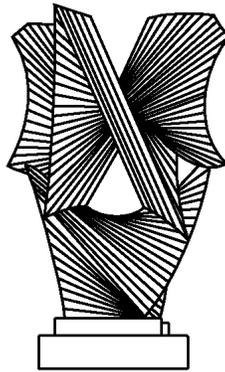


CHICAGO

JOHN M. OLIN LAW & ECONOMICS WORKING PAPER NO. 72
(2D SERIES)



Rethinking Cost-Benefit Analysis

Matthew D. Adler and Eric A. Posner

THE LAW SCHOOL
THE UNIVERSITY OF CHICAGO

This paper can be downloaded without charge at:

The Chicago Working Paper Series Index:
<http://www.law.uchicago.edu/Publications/Working/index.html>

The Social Science Research Network Electronic Paper Collection:
http://papers.ssrn.com/paper.taf?abstract_id=164902

Rethinking Cost-Benefit Analysis

Matthew D. Adler¹ and Eric A. Posner²

INTRODUCTION

The reputation of cost-benefit analysis (“CBA”) among American academics has never been as poor as it is today, while its popularity among agencies in the United States government has never been greater. Many law professors, economists, and philosophers who have written about CBA, believe that it does not produce morally relevant information and should not be used in project evaluation. A few commentators argue that while the information produced by CBA may have some relevance, the procedure itself is deeply flawed. Defenders of CBA form an increasingly beleaguered minority, consisting mostly of applied economists who feel compelled to respond to attacks on the methodological underpinnings of their work. Modern textbooks on CBA are plentiful, and many of them are optimistic about the usefulness of the procedure, but most of them frankly acknowledge its serious flaws and the inadequacy of the standard methods for correcting these flaws.

Government agencies now routinely use CBA. This was not always the case. Before the 1980s, agencies seldom relied on CBA when evaluating regulations and other projects. But executive orders issued by the Reagan and Clinton administrations have since made the use of CBA by agencies common,³ and Congress has enacted numerous statutes requiring agencies to perform cost-benefit analyses.⁴ EPA alone has spent tens of millions of dollars on CBA over the last

¹ Assistant Professor of Law, University of Pennsylvania.

² Professor of Law, University of Chicago. Thanks to Steve Coate, Alon Harel, Saul Levmore, Cass Sunstein, David Weisbach, and participants in workshops at the University of Pennsylvania Law School and the University of Chicago Law School, for helpful comments.

³ EXEC. ORDER NO. 12291, 3 CFR 128 (1981); EXEC. ORDER NO. 12866, 3 CFR 638 (1993). Clinton’s order qualified Reagan’s order a bit but essentially endorsed the use of CBA. See Richard H. Pildes and Cass R. Sunstein, *Reinventing the Regulatory State*, 62 U. CHI. L. REV. 1, 6-7 (1995).

⁴ See Edward R. Morrison, *Judicial Review of Discount Rates Used in Regulatory Cost-Benefit Analysis*, 65 U. CHI. L. REV. 1333, 1333-34 (1998) (and citations therein).

fifteen years.⁵ Other agencies are as committed as EPA to using and improving the techniques of CBA. The academics' skepticism appears to have had no influence on them. What accounts for this divergence between academic opinion and government practice? Are the academic criticisms of CBA valid?

This Article provides a qualified defense of the use of CBA by administrative agencies. It makes the following claims. First, a common criticism of CBA—that it sometimes produces morally unjustified outcomes—overlooks the fact that CBA is a decision procedure, not a moral standard. A decision procedure is a methodology for achieving desirable results, and some decision procedures are more accurate than others. CBA is justified, even if it sometimes produces undesirable outcomes, as long as it produces desirable outcomes more frequently than alternative decision procedures do. We argue that standard alternatives that are proposed in the literature—risk-risk analysis, feasibility-based assessment, direct interpersonal comparisons, and so on—are less accurate than CBA, as long as CBA is used in the right way.

Second, CBA will produce reasonably accurate results only as long as it is used in the right way, and this means that under certain conditions agencies should depart from the traditional understanding of CBA. When a proposed project would affect people who have highly unequal levels of wealth, or who are poorly informed about the consequences of the project, or whose preferences for other reasons fail to register projects that would enhance their well-being, agencies should modify or depart from CBA. The proper modifications will in some cases require the weighting of people's valuations of a benefit of a project; in other cases, the agencies may have to engage in direct comparison of the effect of a project on people's well-being, however hazardous such a judgment might be. The proper adjustments to standard CBA cannot be described at a high level of abstraction, but depend on such things as the competence of agencies, the degree to which they can be monitored by politically responsive actors, and the extent to which people's stated preferences are well-informed and undistorted.

Third, CBA suitably revised to reflect these concerns is consistent

⁵ ECONOMIC ANALYSES AT EPA: ASSESSING REGULATORY IMPACT (Richard D. Morgenstern ed. 1991)

with a broad array of popular theories of the proper role of government. It is commonly and mistakenly believed that CBA presupposes a strict form of utilitarianism that assumes that the government should maximize the satisfaction of people's preferences, even when these preferences are uninformed, distorted, or morally repugnant. By contrast, we argue that CBA, properly understood, is consistent with every political theory that holds that the government should care about the overall well-being of its citizens—and every major political theory, including highly egalitarian theories but excluding extreme libertarianism, holds this view. The use of CBA by agencies in suitable circumstances is consistent with commitments to distributive justice, deontological rights, and other moral values. We also argue that the traditional economic defenses of CBA—based on the Pareto principle and the Kaldor-Hicks principle—are wrong.

We develop this argument in three parts. Part I provides some background on CBA, including a brief history of the procedure and some case studies that show how CBA is typically used by agencies today. Part II describes the mechanics of CBA and explains why the traditional economic defenses of CBA fail. Part III defends CBA on the ground that it is the most suitable decision-procedure for agencies to use in order to maximize the overall well-being of citizens.

I. BACKGROUND

Modern CBA is the outgrowth of three historical developments. The first is the growth of the central government in the United States and other countries over the course of the twentieth century. In the United States the New Deal government initiated the use of CBA in 1936, when Congress ordered agencies to weigh costs and benefits when evaluating projects designed for flood control.⁶ The popularity of CBA among administrative agencies increased rapidly thereafter with the growth of the federal government.⁷ The second development was the rise of Progressivism at the end of the nineteenth century and

⁶ See Ajit K. Dasgupta and D.W. Pearce, *CBA: THEORY AND PRACTICE* 12-13 (1972) (United States Flood Control Act of 1936, which held that projects should be approved if “the benefits to whomsoever they may accrue, are in excess of the estimated costs.”).

⁷ For a brief history, see Theodore M. Porter, *TRUST IN NUMBERS: THE PURSUIT OF OBJECTIVITY IN SCIENCE AND PUBLIC LIFE* 148-89 (1995).

the beginning of the twentieth century. Progressives believed that government can be separated into a realm of value-laden politics, and a realm of administrative expertise which could be based on scientific principles.⁸ The third development was the invention of modern welfare economics, which would supply these scientific principles. Early welfare economists believed that economic concepts could be used to rationalize the implementation of government policies.⁹ Their efforts were encouraged in the 1950s and 1960s when the U.S. government, and the governments of other countries, sought technical assistance in the development of formal procedures of CBA.¹⁰

Modern welfare economics can be traced back to Pareto. Pareto proposed as a principle of evaluation that a project is desirable if it makes at least one person better off without making anyone else worse off.¹¹ The self-evident moral correctness of the Pareto principle has frequently been assumed by economists and others, although as we discuss below it has serious flaws on ethical grounds. But the immediate difficulty posed by the Pareto criterion was that it was too strong. Few projects satisfy the criterion, because just about every worthwhile government project will hurt people, and compensating those people is usually infeasible. This difficulty led to the proposal of hypothetical compensation tests by Kaldor, Hicks, and others.¹² Compensation tests

⁸ See Richard H. Nelson, *The Economics Profession and the Making of Public Policy*, 25 J. ECON. LIT. 49, 52-54 (1987).

⁹ *Id.*

¹⁰ See Dasgupta and Pearce, *supra* note __, at 13. In addition, private companies had independently developed techniques for evaluating their investment decisions, and these techniques could be transferred to the public sector. *Id.*

¹¹ See I.M.D. Little, *A CRITIQUE OF WELFARE ECONOMICS* 84 (2d ed. 1957). One could alternatively date modern welfare economics—or at least the Anglo-American version—to Robbins' attacks in the 1930s on the older "material welfare school," which focused on the material well-being of individuals (as opposed to their utility, in the modern sense), and held that the material well-being of individuals is comparable. See Robert Cooter and Peter Rappoport, *Were the Ordinalists Wrong About Welfare Economics?*, 22 J. ECON. LIT. 507, 520 (1984).

¹² A useful discussion can be found in Little, *supra* note __, at 88-96. See Nicholas Kaldor, *Welfare Propositions and Interpersonal Comparisons of Utility*, 49 ECON. J. 549 (1939); John R. Hicks, *The Foundations of Welfare Economics*, 49 ECON. J. 696 (1939); Harold Hotelling, *The General Welfare in Relation to Problems of Taxation and of Railway and Utility Rates*, 6 ECONOMETRICA 242 (1938).

hold that a project is desirable if its beneficiaries are enriched enough that they could overcompensate those who are hurt by the project. The accomplishment of these tests was to enable the decision-maker to quantify the positive and negative effects of a project and compare these quantities along a single metric. This immensely simplifies the process of evaluating projects, compared to the more-or-less intuitive methods that were presumably used earlier. Moreover, the compensation tests vastly increase the range of projects that can be evaluated, compared to the Pareto test. The compensation tests would be the basis of modern CBA.

The compensation tests, however, were received unenthusiastically by theoretical welfare economists. When the storm of criticism subsided, some economists declared that not only compensation tests, but all of welfare economics, were dead,¹³ a declaration that has been repeated many times since then.¹⁴ Despite these views, CBA obtained a foothold among applied economists and government agencies. Applied economists and agency officials believed that, whatever its problems, CBA was superior to the alternatives. When the government proposed a project, taxpayers and critics demanded a justification, and the most obvious justification was that the project would produce gains that exceeded its costs.¹⁵

Thus, CBA enjoyed a brief period of popularity in the 1960s, despite the absence of a consensus on its theoretical foundation. By the 1970s, however, even applied economists and government agencies had begun to doubt its utility. The emerging problems with CBA were not theoretical, but practical and ideological. As a practical matter, researchers had a great deal of trouble obtaining relevant data, especially for the purpose of valuing environmental resources, human life, and other hard-to-measure goods. The claim that the benefits of a project exceed its costs is not persuasive when the benefits and the costs appear to rely on arbitrary valuations. As an ideological matter, the technical and utilitarian flavor of CBA was unappealing to the political culture

¹³ See J. de V. Graaff, *THEORETICAL WELFARE ECONOMICS* (1957) __.

¹⁴ See, e.g., John S. Chipman and James C. Moore, *The New Welfare Economics 1939-1974*, 19 *INTERN'L ECON. REV.* 547, 548 (1978) (calling modern welfare economics a "failure").

¹⁵ See Dasgupta and Pearce, *supra* note __, at 13.

that reigned during the 1970s.¹⁶ It may be that progress in valuation techniques, and changes in ideology, account for the reemergence of CBA in the 1980s and 1990s—it is too early to tell. The simple truth may be that CBA generally interferes with government regulation, if only by slowing it down, so people oppose CBA when they seek greater regulation (the 1970s) and they favor CBA when they seek less regulation (the 1980s). Whatever the case, the modern rebirth of CBA has not been accompanied by a theoretical defense. The original theoretical objections to CBA have still not been rebutted.¹⁷

Understanding the problems and advantages of modern CBA, however, is difficult. There are two reasons for this difficulty. First, the academic literature on CBA is deeply fragmented, with critics from different disciplines rarely paying attention to each other's arguments. Philosophers object to CBA because they think that it depends on an implausible moral or political theory, like utilitarianism. Economists who object to CBA usually do so on the grounds that it does not allow a complete and consistent ordering of projects, or because it depends on contestable normative premises that cannot be the basis of neutral and scientific advice. Neither group pays much attention to institutional issues, such as the role of agencies in a representative government, yet law professors, who are well positioned to explore these issues, do not say much about them, either. Law professors' discussions of CBA, moreover, frequently overlook the contributions made by economists

¹⁶ See D.W. Pearce and C.A. Nash, *THE SOCIAL APPRAISAL OF PROJECTS: A TEXT ON COST-BENEFIT ANALYSIS* 3-4 (1981); R. Shep Melnick, *The Politics of Benefit-Cost Analysis*, in *VALUING HEALTH RISKS, COSTS, AND BENEFITS FOR ENVIRONMENTAL DECISION MAKING* 23 (P. Brett Hammond and Rob Coppock, eds. 1990) (attributing hostility to CBA to populism and politics).

¹⁷ A large legal literature addresses the related question whether legal rules in general, and the common law in particular, do or should reflect efficiency concerns. See, e.g., *Symposium on Efficiency as a Legal Concern*, 8 *HOFSTRA L. REV.* 485 (1980). This literature is of limited usefulness for our purposes. It does not, for the most part, address CBA, and much of it is concerned about the role of common-law judges. Otherwise, it resembles earlier discussions in the economics literature. More recent treatments in the legal literature include Lewis A. Kornhauser, *Wealth Maximization*, in *THE NEW PALGRAVE DICTIONARY OF ECONOMICS AND THE LAW* (Peter Newman ed. 1998); Mark Kelman, *A GUIDE TO CRITICAL LEGAL STUDIES* 114-50 (1987); Richard S. Markovits, *Duncan's Do Nots: Cost-Benefit Analysis and the Determination of Legal Entitlements*, 36 *STAN. L. REV.* 1169 (1984).

and philosophers.

Second, understanding CBA as described by textbooks is not the same thing as understanding CBA as practiced by government agencies. The economic, philosophical, and legal criticisms are for the most part directed at the textbook version. But for all their enthusiasm for CBA, it is not clear whether agencies use it properly. As we discuss below, agencies sometimes appear to use CBA to rationalize decisions made on other grounds. At other times, agencies may be sincere but they depart from CBA without explaining their departure. For example, they calculate costs or benefits by using national averages when the project would affect a non-representative subset of the population. Or they calculate some of the costs and some of the benefits of the project, while ignoring others. Or they take into account distributional considerations that are external to ordinary CBA and inconsistent with the more ambitious textbook versions of CBA that incorporate distributional weightings. The literature ignores these complications, and instead debates CBA as an abstraction, not as a real practice.

What is the real practice of CBA? It is hard to generalize, but a few examples may help put the problems in context.

*Lead in Drinking Water.*¹⁸ Federal law requires EPA to regulate lead contamination of drinking water. In 1991, EPA decided to revise earlier regulations it had issued under the law, using a CBA of several proposed rules.¹⁹ On the cost side, EPA estimated on an annualized basis, using a 3% discount rate: the cost of treating contaminated water that enters the distribution system, the cost of “corrosion control,” that is, the cost of maintaining appropriate parameters of water quality (pH level, temperature, etc.), the cost of replacing lead pipes, the cost of warning the public of high lead levels and informing it of precautions, and the cost of monitoring water quality. For each rule, EPA calculated total costs by aggregating the cost of implementing the rule in each of the water distribution systems in the United States, which, of course, varied in the severity of lead contamination. The benefits of

¹⁸ Ronnie Levin, *Lead in Drinking Water*, in *ECONOMIC ANALYSES AT EPA*, *supra* note __, at 205.

¹⁹ For a discussion of the earlier rules, see Thomas O. McGarity, *REINVENTING RATIONALITY: THE ROLE OF REGULATORY ANALYSIS IN THE FEDERAL BUREAUCRACY* 29-44 (1991).

the three rules were divided into health benefits for children, pregnant women and fetuses, and adult men. Only some of the benefits were monetized, however. EPA estimated that the cost of medical treatment for children with elevated lead levels would be between about \$300 and \$3200 per child; the cost of compensatory education for children with “cognitive damage” would be about \$5800; and the cost of lost future earnings would be \$4600 per lost IQ point. For adults, EPA estimated a willingness-to-pay of \$1 million to avoid nonfatal heart attacks and strokes, \$628 per case for medical costs and lost productivity because of hypertension, and \$2.5 million per death. Total costs were estimated by multiplying these amounts by the estimated number of cases avoided, and summing the products. Although EPA estimated the cost of the damage of contaminated water to plumbing components, it did not include this estimate in the CBA published with the final rule.

EPA concluded that the total health benefits from corrosion control alone would be \$63.8 billion over a twenty-year period, which vastly exceeded estimated costs of \$4.2 billion. The author of the study (who also worked on the rule) argues that the CBA played an important role in persuading EPA of the hazards posed by lead contamination in drinking water. He also argues that CBA was influential because data were plentiful and the analysis occurred early in the regulatory process.²⁰

*Agricultural Pesticides.*²¹ Federal law authorizes EPA to regulate the labeling and use of pesticides. In 1983 EPA decided to revise earlier regulations, and to evaluate new rules using CBA. On the cost side, EPA determined the costs of workers waiting before entering treated areas, personal protective equipment, notification procedures, training workers in the use of pesticides, decontamination of workers affected by pesticides, emergency assistance and medical care, rule familiarization, and monitoring of the health of selected workers. Because the necessary actions would vary from site to site, different cost estimates were calculated for different kinds of sites, and results were summed. Although the analysis was highly detailed, the final estimate was necessarily rough, because EPA had little information about the

²⁰ *Id.*, pp. 228-29.

²¹ Louis P. True, Jr., *Agricultural Pesticides and Worker Protection*, in ECONOMIC ANALYSES AT EPA, *supra* note ___, at 303.

size of the populations affected by the regulation. EPA did not attempt to attach a monetary value to the benefits of the regulation, stating instead that the regulation it finally chose would reduce the health effects of pesticide use by 80%. These health effects were divided into hospitalized poisonings (300-450 per year), physician-diagnosed but nonhospitalized poisonings (10,000-20,000 per year), undiagnosed poisonings (“a significant number ... very likely to be large” but unquantified), cancer cases (6 or more per year), serious developmental defects (20-52 per year), stillbirths (56-222 per year), persistent neurotoxicity cases (150-300 per year), and others (unquantified). Despite the failure to quantify benefits, EPA concluded that the benefits of the regulation exceeded the costs.

During the rule making USDA argued that in order for the regulation to be cost-justified, it would have to reduce the number of hospitalizations by 239,000, assuming a cost of \$580 each C presumably, an excessive estimate, given that hospitalized poisonings amounted to 300-450 per year. EPA responded that it considered other benefits as well, but it did not quantify them because of deficiencies in the data. One justification for the rule offered by EPA appears to have been that the regulation was not so costly as to cause significant economic disruptions to agriculture, but EPA did not explain what this meant.²² Finally, EPA, according to the author of the study, appeared to take account of distributional considerations without saying so explicitly.²³ The rule would benefit mostly poor agricultural laborers, with its costs being paid by consumers.

These two examples of CBA give rise to several observations. First, EPA, like other government agencies, generally uses a valuation of life or a range of valuations that are invariant across individuals of different wealth, even though CBA will on average attach higher valuations to wealthier people because they can afford to pay more to reduce risk. It is doubtful in the pesticides case that if EPA had quantified benefits, it would have attached a lower valuation of life to migrant farm workers than the national average, even though migrant farm workers are poorer than the average person. Although sometimes EPA

²² *Id.*, at 328.

²³ *Id.*

produces different valuations for different classes of people,²⁴ and although one can properly rely on national averages when projects affect everyone in the nation, EPA clearly violated an important element of CBA.

Second, CBA is frequently hampered by a lack of data or the difficulty of estimating valuations. A striking example is a CBA that attempted to monetize the aesthetic value that people attach to clear air over the Grand Canyon.²⁵ Often, the problem was not that collection of data was infeasible, but that it was precluded by EPA's budgetary and time constraints.

Third, cost-benefit analyses served an important political purpose: by forcing EPA to state clearly the effects of a regulation, they alerted affected groups, which would frequently criticize EPA's estimates. CBA created regulatory transparency.²⁶ There is also some evidence that EPA resisted CBA in order to avoid the political trouble that resulted when affected groups intervened.

Fourth, CBA helped EPA establish priorities. CBA of water contamination by lead revealed that the health costs were significantly higher than the costs produced by more politically salient environmental concerns, such as contamination by radionucleotides.²⁷ Even if one is skeptical about the particular estimates in a CBA, one might use CBA to rank projects by seriousness on the theory that errors wash out.

Finally, when EPA did not use CBA, it was never clear what methodology it relied on. Sometimes, guidance could be found in the statute. But more often, it appears that EPA relied on a kind of implicit CBA. On the one hand, a regulation may appear justified as long as it does not cause too much "economic disruption" to the affected industry. This appears to mean that the regulation will not cause either enormous price increases for consumers or numerous

²⁴ Eloise Trabka Castillo, Mark L. Morris, and Robert S. Raucher, *Great Lakes Water Quality Guidance*, in ECONOMIC ANALYSES AT EPA, *supra* note __, at 419; Leland Deck, *Visibility at the Grand Canyon and the Navajo Generating Station*, in ECONOMIC ANALYSES AT EPA, *supra* note __, at 267.

²⁵ Deck, *supra* note __, at 267.

²⁶ See, e.g., *id.*, at 462.

²⁷ See Nichols, *supra* note __, at 78 ("a week of lead is like a millennium of radionucleotides").

bankruptcies in the affected industry. On the other hand, a regulation may appear justified as long as one can point to fairly concrete health effects, like deaths or cancer cases avoided. However, EPA sometimes issued regulations that caused economic disruption, and often refused to issue regulations even though they avoid more deaths or cancer cases than alternative regulations that were issued. It seems likely that even when EPA refrained from an explicit CBA, it engaged in implicit tradeoffs that were not articulated or quantified, for it is hard to see what else EPA could have been doing.

Hahn analyzed 92 rules issued by five agencies between 1990 and 1995, and found that many of these rules would not pass a cost-benefit test.²⁸ This study is consistent with earlier studies.²⁹ There are a variety of reasons for these results, and one cannot overstate the difficulties of evaluating regulations because of problems in evaluating benefits and determining appropriate discount rates.³⁰ But an important reason appears to be that agencies often do not monetize benefits, and do not engage in an explicit CBA.

In sum, although agencies, like EPA, self-consciously engage in CBA, it is not clear whether their analyses were correctly performed, according to the traditional understanding of CBA.³¹ But criticism of the EPA for deviating from textbook CBA requires a normative theory of CBA. Otherwise, one cannot exclude the possibility that EPA's deviations were normatively justified.

II. THE CONVENTIONAL VIEW OF COST-BENEFIT ANALYSIS

Discussions of CBA are hampered by lack of consistency in the use of terms. The term "cost-benefit analysis" itself is used to refer to

²⁸ Robert W. Hahn, *Regulatory Reform: What Do the Government's Numbers Tell Us?*, in *RISKS, COSTS, AND LIVES SAVED* 208 (Robert W. Hahn ed. 1996).

²⁹ See, e.g., John F. Morall, III, *A Review of the Record*, 10 *REGULATION* 25 (1998).

³⁰ See Lisa Heinzerling, *Regulatory Costs of Mythic Proportions*, 107 *YALE L.J.* 1981 (1998); see also Morrison, *supra* note __, at 1336-37 (discussing inconsistency in discounting across agencies).

³¹ McGarity comes to similar conclusions from his study of regulations issued by EPA, and also by the Department of Agriculture, the Department of Transportation, and the Occupational Safety and Healthy Administration. See McGarity, *supra* note __, at 174-75.

the Kaldor-Hicks standard, to the method of compensating or equivalent variations, and sometimes to any method that requires trading off costs and benefits rather than relying on absolute standards. The critiques and defenses of CBA reflect this confusion. Some people defend CBA because they believe that alternatives do not allow one to make tradeoffs among values, while others criticize CBA because it is inconsistent with the Kaldor-Hicks standard.

This section untangles these problems. It starts with a description of CBA as it is currently understood. We then describe the conventional defenses of it. We should emphasize at the start that most defenders of CBA assume that agencies should maximize the satisfaction of unrestricted preferences, an assumption that is, in our view, both implausible and unnecessary. These defenders, who are usually economists, invite criticism from philosophers who reject the goal of maximizing satisfaction of unrestricted preferences. This debate is unnecessary, because a commitment to CBA does not depend on the goal of maximizing unrestricted preferences.

This issue, however, we defer to Part III. This section discusses the more technical objections to CBA, including the objection that it does not supply a complete and consistent ordering of projects.

A. What is CBA?

1. Measure of Individual Utility Changes

A *project* is any government action, including a law or regulation, that causes a change in the status quo. A project might include the construction of a new highway, the repair of an old bridge, the creation of a national health insurance system, investment in research and development, the enactment of a law against age discrimination—any action that changes the productive capacities of an economy or the distribution of wealth. To evaluate a project, we compare the future “project state of the world” (P) with the “status quo state of the world” (S). In order to avoid biasing one’s decision in favor of the status quo, one should imagine that S and P are both “projects” between which the agency must choose, where the first project involves not changing the status quo. Any benefits from maintaining the status quo, such as minimization of uncertainty, should be treated explicitly as benefits that project S enjoys and project P lacks.

Some people think that a CBA of S and P is conceptually

straightforward, and the only problem posed by CBA is the practical difficulty of collecting data. Suppose that P would be the creation of a new dam. S, the status quo, would mean not constructing the dam. Clearly, a new dam would create benefits: people would enjoy cheaper electricity than under S. Just as clearly, a new dam would be costly in materials and labor that could be used for other projects, and in environmental degradation as well. One might believe that if one can accumulate data on these benefits and costs, the CBA itself would be a simple matter of determining whether the benefits exceed the costs. Unfortunately, matters are not this straightforward.

To understand how CBA works, one must rely on a more precise model of the economy. Consider a two good and two person economy, with goods X and Y, and individuals A and B. P's effect, relative to S, will be to change the amount of Y or the amount of X or both. Usually, a project will not increase both X and Y, but instead increase the amount of one good while reducing the amount of another. Holding the endowments of X and Y constant, if P increases the amount of Y relative to the amount of X, the price of Y (in terms of X) will fall. The individuals, A and B, experience this change in price as an increase in amount of purchasing power for Y and a decrease in the amount of purchasing power for X. Depending on their preferences for X and Y, this change in purchasing power will make one party better off and the other worse off, both better off, or both worse off.

In our example of the dam, P represents the construction of the dam and S represents the decision not to construct the dam. Let Y represent electricity and X represent fish. It is useful to choose one good as the numeraire, by which we mean the baseline good which is used to measure the other goods. If Y is the numeraire, then we talk about measuring fish in terms of electricity. (X could also be the numeraire). But, more generally, we think of the numeraire as representing all the goods in the economy except the other good under consideration, so if Y is the numeraire, then X represents fish, and Y represents everything else, which is denominated in dollars. Then we can measure X in terms of dollars. Although we can thus extend the two-good case to the real economy without causing analytic problems, we will stick to the two-good case, despite its lack of realism, because it is simpler.

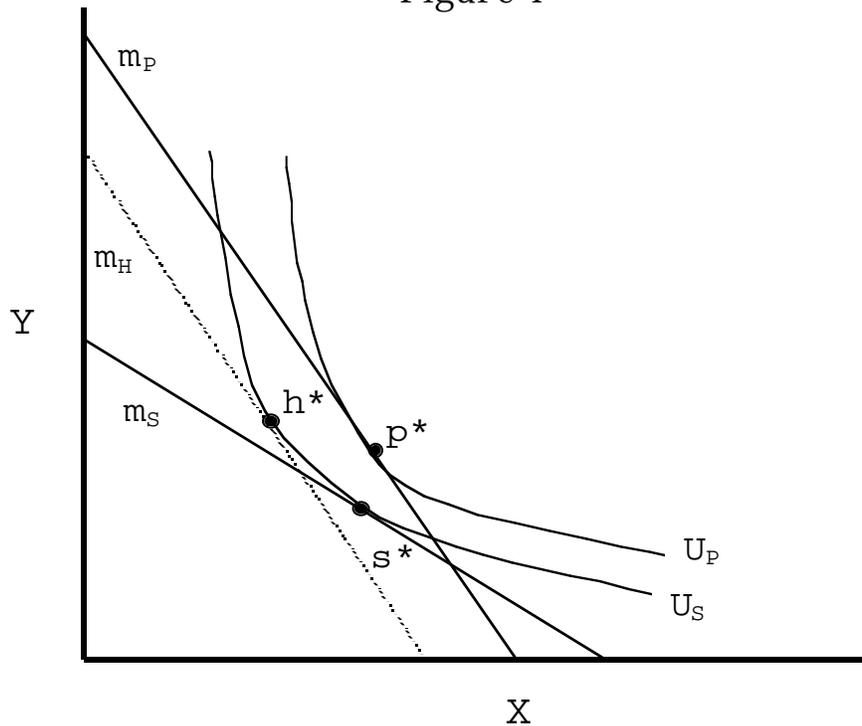
P can have a variety of influences on A and B. The construction of a dam is likely to increase the supply of electricity but reduce the

supply of fish. This means that the construction of the dam will reduce the price of electricity in terms of fish or, what is the same thing, increase the price of fish in terms of electricity. Suppose that in S a person can trade one fish for one unit of electricity, and that in P a person can trade one fish for two units of electricity. If A likes electricity more than fish at the margin, and has a lot of fish and little electricity, P will make A better off. He can exchange the fish for electricity at a higher rate than under S. If A likes fish more than electricity at the margin, and has a lot of electricity and few fish, P will make A worse off. Whereas under S he can trade one unit of electricity for one fish, under P he must use up two units of electricity in order to obtain one more fish. The same comments apply to B. So P can have four effects on the utility of the two people in the economy. It can make both better off relative to S, both worse off, A better off and B worse off, or B better off and A worse off.

Figure 1 illustrates these effects. It shows the effect of the project on a person, say, person B. Under S, B's budget line is represented by m_s , which intersects B's highest indifference curve, u_s , at point s^* . A plausible effect of the dam is to make electricity cheaper and fish more expensive, so if P were implemented, B's budget line would shift to, say, m_p . The steeper slope reflects the fact that electricity is relatively cheaper and fish are relatively more expensive. To see how this works, note that if B does not buy any fish ($X=0$), then B can buy more electricity under P than under S (represented by the fact that m_p intersects the y-axis at a higher point than m_s does). If B does not buy any electricity ($Y=0$), then B can buy fewer fish under P than under S. Under the indifference curves as drawn, P improves B's utility. The project budget line, m_p , intersects a higher indifference curve, u_p (at point p^*). This is just a formal way of showing that the relative decrease in the cost of electricity benefits B more than the increase in the cost of fish. But this is not necessary. If the effect on prices were more extreme (so that m_p intersected m_s above s^*) and B's preferences were different, u_p could be to the left of u_s , so that p^* lay behind u_s . This would be the case if B's relative preference for fish is strong when she is poor, but declines as her endowments increase. In other words, when B is poor, she will exchange a lot of electricity for a few fish; when she becomes wealthier, she will value them more equally, so that she will exchange only equal amounts. At S, she is relatively wealthy. P

increases the cost of fish so much, that she will have to exchange a tremendous amount of the (cheap) electricity in order to satisfy her increased desire for fish, so much that she is worse off than she was under S.

Figure 1



CBA requires that the project's effect on B's utility be converted into units on a metric that enables comparison of the project's effect on B with its effect on other people. An examination of Figure 1 might suggest that we want to measure the distance between s^* and p^* , since these points lie on the curves that represent B's utilities under the two projects. But this distance can be measured only in terms of X and Y: it is the square root of the sum of the squares of the distances traveled along the X and Y axes, as one moves from s^* and p^* . As will become clear shortly, there is no logical connection between this amount and the welfare change for A and B.

One possible solution to this problem is to determine how much

one could take from B in the project state of the world, such that B's utility would be reduced from u_p to u_s . To calculate this amount, one draws a new budget line parallel to m_p and tangent to U_s , which is labeled m_H . Technically, the distance between the points where m_p and m_H intersect the y-axis represents the amount of units of Y that one could take from B in the project world in order to reduce his utility to the level in the status quo. In our example, Y is electricity, so we have converted a utility change into an equivalent change in the amount of electricity that B would consume. At a higher level of abstraction, Y, as the numeraire, represents all goods except X and is measured in dollars. So the distance between the points where m_p and m_H intersect the y-axis is the amount of dollars that would have to be taken from B in the project world in order to reduce his utility to its status quo level. This amount of money is called the *compensating variation* (CV). CBA assumes that B's CV is an adequate representation of the difference in B's utility as between the status quo and the project state of the world.

In our example, B is made better off by P. If B were forced to pay her CV to someone else, then B would fall back to indifference curve u_s . However, B would be at a different point on u_s than under S; she would be at point h^* . B is consuming more electricity and fewer fish than under S, but B is no worse off, because electricity is cheaper and fish are more expensive, thanks to the change in prices caused by P. If P reduced B to a lower indifference curve, then the CV would be negative, and represent the amount of money that would be necessary to raise B's utility to that of S.

Before turning to the question of how to aggregate the CVs of multiple parties, we should point out a difficulty with the analysis so far. In our two-good example we chose to measure the utility effect of the project by using prices in the project state of the world rather than prices in the status quo, but we could have taken the opposite approach. Under the method of *equivalent variation* (EV), one asks how much money one must give (or take from) the individual in the status quo, in order to raise (or lower) his utility level to that of the project state of the world. Graphically, one draws the hypothetical budget line, m_H , parallel to m_s and tangent to the indifference curve U_p , rather than parallel to m_p and tangent to U_s . The distance between the points at which m_s and m_H intersect the y-axis represents the EV. There is no theoretical reason to prefer CV over EV, and the choice between them creates

some indeterminacy in CBA. However, defenders of CBA point out that the data used to calculate CVs and EVs are so crude that, as a practical matter, the two measures will produce similar results.

The difference between CV and EV results from income effects. As a person obtains more goods, his relative preferences among goods may change. For example, a relatively poor person who has few fish and few units of electricity may initially be indifferent between obtaining one more unit of one good and one more unit of the other, because he values an additional unit of each good equally; but as he obtains more and more units of both goods, he may begin to prefer electricity, which can be used to power his television set, to fish, for which he no longer has an appetite. As income increases, a person will not necessarily want to continue to consume two goods at the same rate. It seems reasonable to assume that preferences for most, if not all, goods follow this pattern, and therefore that CV and EV will almost always diverge. Thus, cost-benefit analyses that rely on CV may produce different results from analyses that rely on EV.³²

2. Aggregation

The reason for determining B's CV is to enable a comparison of the effect of P and S on B and A. Recall that in our example P placed B on a higher indifference curve. Now P might also place A on a higher indifference curve, in which case P is Pareto superior to S. We will discuss Pareto superiority below. For now it is sufficient to note that Pareto superiority may be a sufficient condition for approving a project, but few, if any, actual projects are Pareto superior to the status quo, and the reason for using CBA is that the Pareto standard cannot be used to evaluate the vast majority of government projects. For the purpose of example, we must suppose that P injures someone, and since it benefits B, we assume that it hurts A. Because A, a fisherman, loves fish more than electricity, P's effect of increasing the price of fish in terms of electricity reduces the extent to which A can satisfy his preferences. A's CV is calculated in the same way that B's CV is

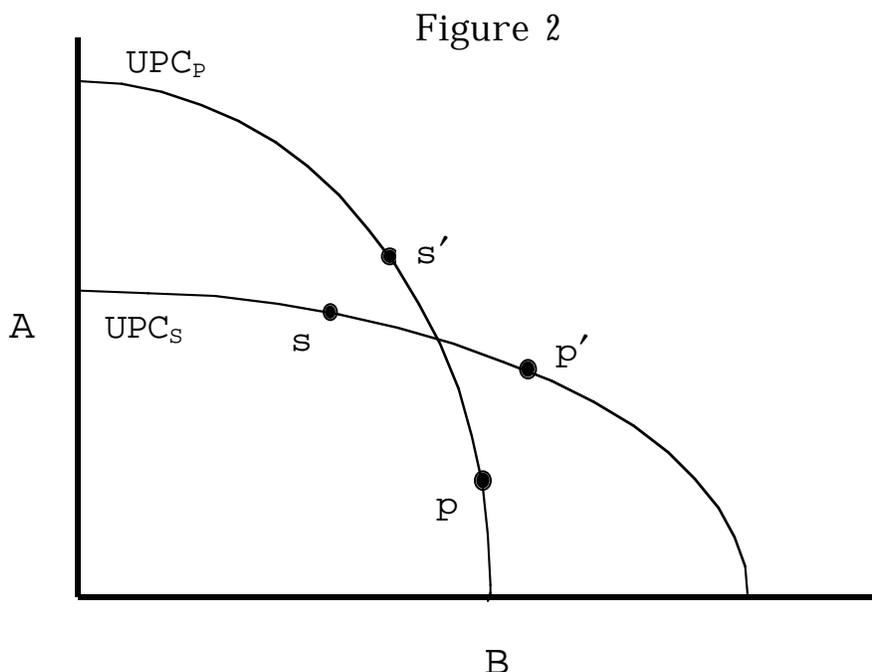
³² An additional problem is that people may, for poorly understood psychological reasons, be willing to pay more for a good than they are willing to accept to give up the same good, holding everything else equal. (CV and EV are measures of WTP if positive; and WTA if negative.) See Elizabeth Hoffman and Matthew L. Spitzer, *Willingness to Pay Vs. Willingness to Accept: Legal and Economic Implications*, 71 WASH. U.L.Q. 59 (1993).

calculated, but this time CV is a negative number. We would have to give A money in the project state of the world in order to make him as well off as he was in the status quo, unlike B, from whom we had to take money. Then we approve the project if the sum of A's CV and B's CV exceeds zero; otherwise, the project is rejected.

Figure 2 depicts this analysis. The two curves are utility possibility curves (UPC). They measure the amounts of utility that A and B can jointly obtain under different technologies. Given a particular production technology, one can give all the goods either to A or B (represented by the intersections of the curve at the axes) or one can split the goods between them. UPC_s represents possible utility distributions in the status quo, and UPC_p represents utility distributions in the project state. Points s and p represent the distributions of utility in the status quo and under the project. If p were in the quadrant northeast of s (for example, where s' is), then p would be Pareto superior to s . Both parties would have higher utility under P than under S . In our example, however, p must be below and to the right of s , in order to represent the fact that the project makes B better off and A worse off. The question is whether the project makes B *sufficiently* better off, relative to A. CBA tells us that p does make B better off by an amount greater than the amount by which it makes A worse off. This can be shown graphically. If the state implements the project and it can engage in costless lump-sum transfers, it can move the utility distribution from the status quo (s) to the project state of the world (p), and thence along UPC_p to a hypothetical world (p') which is Pareto superior to the status quo. This procedure shows that after the hypothetical transfer, both A and B are better off than in the status quo.

One should be clear about what one shows by aggregating CVs. One does not show in a straightforward way that B's well-being is enhanced more than A's well-being is reduced. Rather, one shows that B satisfies her preferences to a greater extent under P than under S , that A satisfies his preferences to a smaller extent under P than under S , and that B's improvement is such that B could more than compensate A for his loss. One reason for this result may be that B's preference for the electricity made cheaper by P is more *intense* than B's preference for the good made more expensive by P . B prefers electricity much more than fish; A is close to indifferent. Thus, if B gave A some extra units of electricity, A would be compensated for his loss of fish

and B would still be better off. Holding everything else equal, this difference in intensity of preferences—so long as unrestricted preferences are entitled to respect³³—may justify a project that makes electricity a lot cheaper and fish a little more expensive.



But another reason for the result may be that as B accumulates more fish and electricity, her relative preference for electricity increases—her stomach being full of fish, she wants to watch more television—whereas A, at a low endowment, is indifferent between fish and electricity but still highly desirous of both. A needs fish for food and electricity for warmth, and wants them intensely, but also wants them equally, so he is not willing to give up a lot of one good in order to obtain a little of the other. One's CV reflects not just the relative intensities of one's preferences, but how these preferences

³³ This is a controversial proposition, which we discuss in Part III.

change as one's endowments increase. And yet this strikes a false chord. If B has plenty of fish but a modest appetite, she may be willing to trade lots of fish at the margin for a little electricity. So she is willing to pay a lot for P. Meanwhile, A does not have very many fish or much electricity, and values fish slightly more than electricity, but still values both a great deal. P makes him worse off, because he must reduce his consumption of fish, and the increased warmth does not offset that loss. But he is not willing to pay much (in terms of electricity units) to avoid P, because at his low endowment electricity matters as much to him as fish do. At a more abstract level, wealthier people have higher CVs for projects than poorer people do, because they value their last dollar less than poorer people do. Millionaires do not value \$100 as much as poor people do; therefore, millionaires would be willing to pay \$100 for a project that increases their utility slightly, and poor people would not be willing to pay \$100 to prevent the same project, even though that project may reduce their utility quite a bit.³⁴

Some scholars argue that this bias in favor of wealthy people is a decisive objection to CBA.³⁵ One response is that if CBA benefits wealthier people more than poor people, but at the same time makes wealthier people better off by more than it makes poor people worse off, the bias can be reversed through redistribution of wealth, in which case enough people will be better off as a result. This response correctly points out that an undesired consequence of CBA can be remedied, but it does not deal with the deeper philosophical difficulty, which is that CBA may not measure anything that we care about. CBA reflects both preference intensity, which we do care about, and wealth, which we do not care about; but can these influences be untangled?

Before answering this question, we should point out some practical consequences of this philosophical difficulty. Suppose that our two people, A and B, have the following endowments in S and P, of the two goods, X and Y:

³⁴ It is not the case that wealthier people will *always* have higher CVs for a given project than poor people do. To give a simple example, a project that generates one dollar for a wealthy person and takes one dollar from a poor person, will produce equal CVs for the wealthy person and the poor person—one dollar. The bias discussed in the text refers to a tendency over a large number of different projects.

³⁵ See C. Edwin Baker, *The Ideology of the Economic Analysis of Law*, 5 PHILOSOPHY & PUB. AFF. 3 (1975).

	A	B
S	2, 0	0, 1
P	1, 0	0, 2

If X refers to fish, and Y refers to units of electricity, then the project of building a dam can be seen as reducing the number of fish from 2 to 1 and increasing the units of electricity from 1 to 2. B is a steel mill owner who benefits from cheap electricity if the dam is constructed and A is a fisherman who benefits if the dam is not constructed. We make the plausible assumptions that for a particular good each party prefers more of that good to less of that good, and that each party prefers equal amounts of each good to unequal amounts of each good, that is, that each party obtains diminishing marginal utility from consumption of a good.

An agency must choose between S and P. If P is implemented, then B obtains an extra unit of electricity. If this unit were costlessly transferred to A, then A's utility would exceed his status quo utility, because A prefers (1,1) to (2,0). Meanwhile, B would be no worse off, with (0,1), than in the status quo. Therefore, P is superior to S. However, if S is implemented, then A obtains an extra unit of fish. If this unit were costlessly transferred to B, then B's utility would exceed her status quo utility, because B prefers (1,1) to (0,2). Meanwhile, A would be no worse off, with (1,0), than in the status quo. Therefore, S is superior to P. Accordingly, the agency has no grounds for preferring S or P. This problem is called the "Scitovsky paradox."³⁶

The Scitovsky paradox is illustrated in Figure 2. As we saw above, P defeats S, because the costless redistribution of p, p', is northeast of s. But S also defeats P, because the costless redistribution of s, s', is northeast of p. As long as the utility possibility curves cross, this indeterminacy is possible. Utility curves do not cross when preferences are identical and homothetic, which means that a person's relative demand for a good does not change with income. It is clear, however,

³⁶ See Tibor Scitovsky, *A Note on Welfare Propositions in Economics* 9 REV. ECON. STUD. 77 (1941); our example is taken from Yew-Kwang Ng, WELFARE ECONOMICS: INTRODUCTION AND DEVELOPMENT OF BASIC CONCEPTS 59-66 (1980). See also Robin Boadway and Neil Bruce, WELFARE ECONOMICS 96-101 (1984).

that people's preferences are not identical and homothetic.³⁷

To understand the challenge posed by the Scitovsky paradox, one must recall that the original purpose of CBA and related concepts of economic efficiency was to separate problems of distribution and problems of wealth or welfare improvement. Economists hoped to distinguish efficiency questions, which their expertise qualified them to address, and distributive questions, to which economic learning had nothing to contribute. The claim was not that distributional questions were unimportant, and that redistribution of wealth was unjustified; it was rather that the economist had nothing useful to say about how wealth should be distributed.³⁸

Thus, it has never been an objection to efficiency standards that they cannot evaluate purely redistributive projects. For example, suppose the government proposes a project that would change $\{(1,0), (3,3)\}$ for A and B, to $\{(2,2), (2,1)\}$. This project does not pass a cost-benefit test, because A cannot *overcompensate* B from his gain; but this only means that the project cannot be justified on efficiency grounds but only (if at all) on distributive grounds. The economist does not condemn this project; he or she expresses no opinion about it. A line of thought does hold that when the government wants to redistribute wealth, the most cost-effective way of doing so is through taxes and transfers, rather than the building of dams in poor areas or toxic waste dumps in rich areas.³⁹ But the choice of how much wealth to redistribute is outside the economist's area of expertise.

We are less interested in the role of the economist as in the role of the government agency, but an analogous argument holds. When the Environmental Protection Agency or the Food and Drug Administration decides whether to approve a project, it might seem that it should do so on the basis of overall well-being (however defined), and not as a way to redistribute wealth from one segment of the population to another. If wealth should be redistributed, independent efforts to do so by uncoordinated agencies seem less likely to succeed than adjustment of taxes and welfare benefits by Congress.

³⁷ Indeed, UPCs may cross multiple times, which can happen under plausible assumptions about preferences. See Little, *supra* note __, at 107-08.

³⁸ See *supra* note __.

³⁹ See Louis Kaplow and Steven Shavell, *Why the Legal System Is Less Efficient than the Income Tax in Redistributing Income*, 23 J. LEGAL STUD. 667 (1994).

The purpose of CBA is to separate out the distributional issue and isolate the efficiency issue, so that the agency will evaluate projects solely on the basis of their efficiency.⁴⁰

But there remains the question whether such a separation is possible. The Scitovsky Paradox arises because the efficiency effect of a project is not independent of its distributional effect. When a project has a sufficient effect on a person's wealth, it will change how much that person values some goods relative to other goods, which means that in the project world he or she will have a new CV for going back to the status quo. This new CV will not be the same as the CV for going from the status quo to the project world. If the other party's CV changes in the right way, the reversal will occur. The reversal will not *necessarily* occur. It will not occur if, for example, a project's effect on people's endowments is small, or the people affected by the project have similar endowments.⁴¹ But even if the reversal will not occur, its possibility haunts the entire project of CBA, because it shows that people's valuations depend on their relative wealth as well as on the intensity of their preferences.⁴² If you care about overall welfare in the ordinary meaning of that term, you should care about satisfying intense preferences, but there is no reason to care about satisfying the preferences of wealthier people more than the preferences of poorer people.

When an agency makes a decision, its decision will have definite distributional effects even assuming that Congress could later make costless lump-sum transfers. If the economy is at point *s*, implementation of the project forces Congress to choose a distribution

⁴⁰ This is a common view. See, e.g., Kaldor, *supra* note __; A. Harberger, *Three Basic Postulates for Applied Welfare Economics: An Interpretative Essay*, 9 J. Econ. Lit. 785 (1971); E. Mishan, *COST-BENEFIT ANALYSIS* (1992).

⁴¹ Cf. Neil Bruce and R.G. Harris, *Cost-Benefit Criteria and the Compensation Principle in Evaluating Small Projects*, 90 J. POL. ECON. 755 (1982).

⁴² See Blackorby and Donaldson, *supra* note __, at 490-91, for a clear statement of this point. Some economists simply argue that a project should not be approved when approval would lead to a Scitovsky paradox. See, e.g., Scitovsky, *supra* note __; Little, *supra* note __, at 112 (who also requires a distributive criterion). This approach does not remove all difficulties: it can produce intransitivities, where project A dominates B, B dominates C, and C dominates A. See Boadway and Bruce, *supra* note __, at 99-100. As the test becomes more rigorous, it produces fewer such inconsistencies, but also enables fewer comparisons of projects.

along UPC_p : for example, p' , p (if Congress favors B), or anywhere else along this curve. However, Congress may believe that the optimal distribution is at s' or any nearby point. By implementing the project, the agency prevents Congress from reaching s' , whereas if it did not implement the project, Congress would be able to reach s' . The agency, by hypothesis, has no grounds for preferring p' over s' ; therefore, it should not take an action that affects Congress' ability to reach s' or p' . That is, it should not, if it is not permitted to make distributive judgments. But the point is that the agency must make *some* decision, and since its decision affects distribution, it should take distributional considerations into account. So the whole project of enabling agencies to make efficiency judgments, while permitting them to avoid distributional judgments, is impossible.⁴³

B. The Conventional Defenses of Cost-Benefit Analysis

Economists have defended CBA in several ways. These defenses all have in common an implicit commitment to the view that people's unrestricted preferences should be respected.⁴⁴ We should make clear that we reject this view. However, because the view is important, and the conventional defenses contain several influential ideas, we will spend this section describing and criticizing them. We classify these defenses as the Pareto defense, the Kaldor-Hicks defense, and the utilitarian defense.

⁴³ For the sake of brevity, we do not address several other objections to cost-benefit that have been influential in the literature. These objections include criticisms of CBA's treatment of (i) risk, (ii) the discounting of future benefits, and (iii) valuation of life and other hard-to-measure goods. For discussions, *see, e.g.*, Herman B. Leonard and Richard J. Zeckhauser, *Cost-Benefit Analysis Applied to Risks: Its Philosophy and Legitimacy*, in *VALUES AT RISK* (Douglas MacLean ed. 1986) (on risk); Richard Layard and Stephen Glaister, *Introduction*, in *COST-BENEFIT ANALYSIS* (Richard Layard and Stephen Glaister eds. 1994) (on discounting); and W. Kip Viscusi, *FATAL TRADEOFFS: PUBLIC AND PRIVATE RESPONSIBILITIES FOR RISK* (1992) (valuation of life).

⁴⁴ Notable exceptions include Harsanyi, *see* J.C. Harsanyi, *Game and Decision Theoretic Models in Ethics*, in *1 HANDBOOK OF GAME THEORY WITH ECONOMIC APPLICATIONS 703-04* (Robert J. Aumann and Sergui Hart eds. 1992), and Sen, *see* Amartya Sen, *COLLECTIVE CHOICE AND INDIVIDUAL WELFARE* ch. 4 (1982). Although both scholars appear to support CBA, at least under certain conditions, neither of them has attempted to reconcile his views with the traditional approach to CBA.

1. *The Pareto Defense*

A standard defense of CBA is that it provides a sufficient approximation of Pareto superior projects. This argument naturally assumes that the Pareto standard is ethically desirable. Let us begin with that claim.

At first sight, the Pareto standard appears normatively attractive. A project that harms no one, and makes at least one person better off, is consistent with a wide variety of moral commitments, including classical liberalism and utilitarianism. It seems to be consistent with common-sense morality.

Standard objections to the Pareto principle include the obvious point that people may make incorrect choices, so their indifference curves will not describe distributions that are systematically related to their actual welfare. A drug addict may reach a higher indifference curve as a result of a project that reduces the price of drugs, but most people would condemn such a project rather than support it. Another objection is that a Pareto superior allocation may be distributively unjust. A project that generates \$1000 for a rich person and nothing for a poor person aggravates wealth inequality. A third objection is that the Pareto standard assumes a commitment to ethical individualism, with the satisfaction of preferences taking priority over the enhancement of community values. But, it is claimed, ethical individualism does not accord with our moral intuitions. The standard responses to these claims are that agencies should sometimes ignore distorted preferences or that preferences are not usually distorted; the government can redistribute wealth in order to achieve distributive justice; and ethical individualism *does so* accord with our moral intuitions. We address this debate in Part III.

More significant, for our purposes, is the problem that the Pareto standard cannot supply a sufficiently complete ordering of projects. It is likely that the Pareto standard would reject desirable projects that would be approved under an uncontroversial social welfare function. For example, a vaccine that improved the health of millions of people but required a tax of \$1 on someone unaffected (who is not altruistic) would violate the Pareto standard but be required by any plausible moral theory. (Suppose that the unaffected person could not be excused from the tax except at great administrative cost resulting in such high taxes for everyone else that the vaccine would no longer improve their well-

being.) All utility-enhancing government projects probably violate the Pareto standard. Although one might argue that Pareto superiority could be a sufficient condition for a project, we doubt that this claim is of any importance, since the Pareto standard is almost completely indeterminate.

This is where CBA comes in. Defenders of CBA argue that CBA provides a useful approximation of the Pareto standard, while also allowing a *more* complete ordering (though, as we saw, not a fully complete ordering). People have proposed two arguments that CBA approximates the Pareto standard.

The first argument is that although CBA does not require the losers to be compensated, the government can (and should) compensate the losers by taxing the winners after the project is implemented. The problem with this argument is that if the government can and does tax the winners and compensate the losers, the project is Pareto superior to the status quo and CBA is not necessary. But, as we saw, very few projects are truly Pareto superior, because the administrative costs of identifying everyone injured by a project and transferring money to those people from the winners would overwhelm the project's benefits.

The second argument is that people will expect in the aggregate, over time, to be benefited by projects as often as they are injured by projects.⁴⁵ A government that uses CBA will thus, in general, have "wealthier" citizens than a government that fails to use CBA. Suppose we must decide today whether for now on the government should use CBA. If CBA increases the wealth of everyone in the aggregate, then every person expects *ex ante* to be better off with CBA than without it. In this *ex ante* sense, CBA is Pareto-superior to the status quo. The problem with this argument is that in a many-good economy, CBA will tend to favor people who have a low opportunity cost for money and intense preferences. There is no reason to believe that the people who usually injured by projects are usually the same as the people who are usually benefited by projects. And although the government might redistribute wealth through the welfare system, there is no reason to believe that the beneficiaries of welfare are the same as the people

⁴⁵ See A. Mitchell Polinsky, *Probabilistic Compensation Criteria*, 86 Q.J. ECON. 407 (1972), for an attempt to formalize this idea, which goes back to Hicks, *supra* note ____.

injured by projects implemented by the government. Thus, the Pareto standard will not be satisfied.

Another line of thought accepts the force of these criticisms and argues that CVs should be calculated using distributional weights. One possibility is to multiply a person's CV by his marginal utility of money. Because poor people have higher marginal utilities of money than do rich people, the weighting system will inflate their CVs relative to those of rich people, with the result that approval of projects would not be biased in favor of rich people. The problem here is that of determining people's marginal utilities of money. Most economists appear to believe that the difficulties that would be involved are insurmountable, and so this scheme and related schemes have not had much influence.⁴⁶ Another possibility is to give the agency the power to make distributive judgments, and forbid it to approve projects that would worsen the existing distribution of wealth.⁴⁷ But this standard would be too restrictive, preventing the approval of projects that may have detrimental but tiny effects on the distribution of wealth while helping a great many of people in a significant way. Because the literature in this area is inconclusive and has not had much influence on government policy, we will not pursue this line of argument.

2. *The Kaldor-Hicks Defense*

Some scholars defend CBA on the grounds that it approximates the Kaldor-Hicks standard. The Kaldor-Hicks standard states that a project is desirable if it makes the winners better off by an amount sufficient to overcompensate the losers, if the losers could be compensated through a costless lump-sum transfer. More precisely, state P Kaldor-Hicks dominates state S if it is possible to (costlessly) redistribute goods in state P so as to produce a distribution that is Pareto superior to the distribution in state s. Looking at Figure 2, note that Kaldor-Hicks ranks Project p over Pareto noncomparable Project s. To see why, observe that a costless redistribution of the bundle of goods represented by p would allow a move to p', which is Pareto-superior to s.⁴⁸ This defense assumes that the Kaldor-Hicks standard

⁴⁶ See Little, *supra* note __, at 120-28.

⁴⁷ Little proposes such a criterion, although he is not speaking specifically about agencies. See *id.*, at 112.

⁴⁸ We avoid two complications. First, the Kaldor-Hicks standard actually refers

is normatively defensible. The difference between CBA and Kaldor-Hicks is that CBA uses money as the numeraire, whereas Kaldor-Hicks, a more general criterion, does not use a numeraire.

Most economists appear to concede that the Kaldor-Hicks standard is not, by itself, normatively desirable.⁴⁹ The problem with the Kaldor-Hicks standard is that hypothetical compensation is not real compensation. The loser when a project is approved is not consoled by his compensation in a hypothetical world; the Kaldor-Hicks standard lacks precisely what makes the Pareto standard attractive.

So the Kaldor-Hicks standard is usually defended by reference to the Pareto standard. Indeed, it is often called the “potential Pareto” standard. The argument is that although an individual might lose as a result of one project, he or she is also likely to win as the result of another project, so over time the gains and losses will even out, and everyone (or almost everyone) will be better off if the Kaldor-Hicks standard is used than if some alternative were used. In addition, distributive problems can be solved with the tax and welfare system. But this argument is no different from the claim that CBA approximates the Pareto standard, an argument we rejected in the prior section. As noted above, the only difference between using CVs and using the Kaldor-Hicks criterion is that the former uses money as a numeraire, whereas the latter does not use a numeraire; but this difference does not affect the conclusion that the standards are biased in favor of people who are wealthier (in money or goods). Moreover, justifying CBA on the basis of Kaldor-Hicks is harder than justifying Kaldor-Hicks on the basis of the Pareto principle, because, as Boadway has shown, CBA is a necessary but not sufficient condition for Kaldor-Hicks.⁵⁰ But this argument, and the literature it has spawned, is not of any practical

loosely to two separate standards, the Kaldor standard and the Hicks standard. The choice between Kaldor and Hicks parallels the choice between CV and EV, and introduces further indeterminacy. Second, we ignore the debate about whether the hypothetical transfer should be considered costless or should be understood to require the costly redistributive instruments at the government's disposal—a debate which, in our view, is idle.

⁴⁹ See, e.g., Charles Blackorby and David Donaldson, *A Review Article: The Case Against the Use of the Sum of Compensating Variations in Cost-Benefit Analysis*, 23 CANADIAN J. ECON. 471, 472 (1990).

⁵⁰ Robin W. Boadway, *The Welfare Foundations of Cost-Benefit Analysis*, 84 ECON. J. 426 (1974).

interest. Because Kaldor-Hicks is, taken as a moral principle, indefensible, CBA cannot be justified by reference to Kaldor-Hicks. To defend CBA, one must appeal to some other moral principle.⁵¹

3. *The (Unrestricted) Utilitarian Defense*

A final defense of CBA is that CBA is justified assuming a commitment to utilitarianism with unrestricted preferences. We do not know of any sustained defenses of this position, but it seems to be implicit in the work of some authors.⁵² We also do not think that this version of utilitarianism is plausible. But assuming it were, how could CBA be defended on the ground of unrestricted utilitarianism?

Initially, we must define what we mean by utilitarianism. Classical utilitarians like Bentham and Mill used the concept of utility to refer to a distinct mental state, something like a feeling of happiness or well-being. Maximizing aggregate utility, then, meant increasing the happiness of as many people as possible. In principle, the amount of happiness that a person has could be quantified. For reasons that we discuss in Part III, most modern economists reject this view. Modern economists hold that utility refers to the extent that a person satisfies his or her (unrestricted) preferences. A utility function ranks states of the worlds according to the extent to which a person satisfies these preferences. Thus, modern utility is *ordinal* rather than *cardinal*. This means that if by reducing the price of fish, we increase a person's utility function or (what is generally the same thing) raise that person to a higher indifference curve, we cannot say how much better off is that person, but only that he or she is better off. Indeed, to be more precise, we can say only that the person can satisfy his or her (unrestricted)

⁵¹ We do not pursue here an argument that Kaldor-Hicks is justified on the basis of hypothetical consent, see Richard A. Posner, *The Ethical and Political Basis of the Efficiency Norm in Common Law Adjudication*, 8 HOFSTRA L. REV. 487 (1980), and, for criticisms, see, e.g., Jules Coleman, *Efficiency, Utility, and Wealth Maximization*, 8 HOFSTRA L. REV. 509 (1980); or that it is consistent with people's moral intuitions. See Richard A. Posner, *Utilitarianism, Economics, and Legal Theory*, 8 J. LEGAL STUD. 103 (1979), and, for criticisms, see, e.g., Ronald Dworkin, *Is Wealth a Value?*, 9 J. LEGAL STUD. 191 (1980). Other contributions to this debate can be found in *Symposium on Efficiency as a Legal Concern*, *supra* note ____.

⁵² See Pearce and Nash, *supra* note ____, at pp. 26-27. Pearce and Nash do not themselves accept this approach. *Id.*

preferences to a greater degree. But even if we assume that we, as utilitarians, want to do this for as many people as possible, we run into a problem. The problem is there is no non-arbitrary method for ranking social states on the basis of their effect on the ordinal utility functions of all people.⁵³ For any given project, some people will prefer it to other projects and other people will prefer the other projects to the first project. An agency that obtained sincere answers to questions about people's ranking of the projects would not be able to use that information to rank the projects in a non-arbitrary way.

This result is best understood to be a problem about information.⁵⁴ Welfare economists like to assume ordinal utilities, rather than cardinal utilities, because the former place much less demand on the capacity of the decision-maker to obtain information. To compare cardinal utilities, the decision-maker must have some idea about how much "happiness" or "welfare" a person would experience under alternative states of the world, and be able to compare this level of utility with the utilities of other people. To compare ordinal utilities, the decision-maker must have some idea about whether a project increases one person's satisfaction of his preferences by "more" than it reduces some other person's satisfaction of preferences. Some economists controversially have argued that such interpersonal comparisons of welfare are impossible.⁵⁵ More plausibly, they are just very so hard.⁵⁶ They are so hard that an agency that insisted that a project made some people happier by an amount greater than the amount of happiness lost by other people would have a hard time persuading anyone of the basis for its conclusions.

Still, one can make sense of the idea that agencies should use

⁵³ See Kenneth Arrow, SOCIAL CHOICE AND INDIVIDUAL VALUES (1951).

⁵⁴ Arrow's theorem can be evaded by assuming the possibility of interpersonal comparability. See Anandarup Ray, COST-BENEFIT ANALYSIS: ISSUES AND METHODOLOGIES 33 (1984).

⁵⁵ See Lionel Robbins, *Interpersonal Comparisons of Utility: A Comment*, 43 ECON. JOURNAL 635 (1938). For a discussion and criticism of Robbins' position, see Little, *supra* note __, at 50-66. A recent article states that "most economists think interpersonal comparisons are nonsense." Robert A. Pollak, *Welfare Comparisons and Situation Comparisons* 50 J. ECONOMETRICS 31 (1991). However, the author, who is himself agnostic on the matter, *id.*, at 39, does not supply citations, and we have not found recently published work that makes such a strong claim.

⁵⁶ See Pollak, *supra* note __, at 37-43, for a discussion of the difficulties.

CBA in order to maximize utility. Suppose that every person in society has identical endowments. Then it can plausibly be assumed that each values an extra dollar by the same amount. People's CVs would reflect their utilities exactly, and a project that passed a CBA would also increase aggregate utility. Now suppose that people's endowments are not identical. In order to aggregate utility, one could not use CVs, because they are distorted by the differences in endowments. A rich person might be willing to pay more for a project than a poor person is, yet it is likely that his marginal utility of money is less than the poor person's, so approval of the project would result in a reduction in aggregate utility. In principle, one could weight CVs in order to eliminate this distortion, but in order to determine the proper weighting system, an agency would have to determine everyone's marginal utilities of money. Most economists appear to believe that such a determination would be too difficult.⁵⁷ So economists are faced with a Hobson's choice. One branch of the literature proposes that economists should evaluate projects on the basis of social welfare functions that include proper distributional weightings. Utility is assumed to be cardinal under this approach: different people's utilities can be weighted and summed. This approach is, in a sense, intellectually rigorous, but it is not useful because it is too demanding on the decision-maker, and one does not observe agencies (or anyone else, as far as we know) using such ambitious social welfare functions in the real world. The other branch of the literature holds that economists should evaluate projects on the basis of unweighted CVs.⁵⁸ This approach, for which ordinal utilities are sufficient, is perhaps less intellectually respectable, but has had more influence. The pure approach is impractical; the practical approach is impure. It is relatively straightforward to aggregate and compare CVs, but the outcome does not necessarily tell one whether a project enhances welfare, whether understood ordinally or cardinally.⁵⁹

⁵⁷ See Pearce and Nash, *supra* note __, at 27. A powerful critique can be found in Little, *supra* note __, at 120-27.

⁵⁸ That is, the social welfare function does not include distributional weightings.

⁵⁹ Most textbooks on CBA recommend distributive weighting and discuss various methods. See, e.g., Boadway and Bruce, *supra* note __, at 271-91; D.W. Pearce and C.A. Nash, *THE SOCIAL APPRAISAL OF PROJECTS* 31-37 (1981); Richard O. Zerbe, Jr., and Dwight D. Dively, *BENEFIT-COST ANALYSIS IN THEORY*

Finally, it is worth mentioning an argument that, we believe, has currency among economists although it is rarely defended in print. This argument is that CBA is desirable because there are no superior alternatives that provide determinate, or relatively determinate, prescriptions.⁶⁰ This argument assumes that if agencies engaged in some sort of direction utilitarian regulation, they would be unlikely to evaluate projects in a consistent way. If CBA provides only a feeble approximation of utilitarianism, that is better than no guidance at all.

This argument might seem odd. Why would a poor guide be better than no guide at all? If someone proposed a method that required the approval of all projects whose titles have at least 20 letters and the disapproval of all projects whose titles have fewer than 20 letters, the method would produce determinate results but not desirable ones. One way to understand this argument is to imagine that a person is lost in the woods. Having no theory about how to get out of the woods, he or she walks around randomly. A method that does not tell the person how to get out of the woods may still be helpful. It is

AND PRACTICE 236-53 (1994). However, they do not show that the weighting systems are practical, and indeed many textbooks express doubts about their practicality. See, e.g., Mishan *supra* note __; Anandarup Ray, COST-BENEFIT ANALYSIS: ISSUES AND METHODOLOGIES 22-31 (1984); Robert Sugden and Alan Williams, THE PRINCIPLES OF PRACTICAL COST-BENEFIT ANALYSIS 206-07 (1978). So the literature leaves one in doubt about what a proper CBA entails. As mentioned in the text, agencies do not appear to use explicit distributional weightings of the sort recommended by textbooks, although they may engage in such weightings surreptitiously or informally. See the discussion of pesticide regulation, *supra* note __, and accompanying text. Some scholars argue that the economist should not use distributive weights but should disaggregate the costs and benefits of a project for particular groups, and allow the policy maker decide whether its distributional consequences are acceptable. See, e.g., A.R. Prest and R. Turvey, *Cost-Benefit Analysis: A Survey*, 75 ECON. J. 683, 701-02 (1965).

⁶⁰ See Alan Williams, *Cost-Benefit Analysis: Bastard Science? And/or Insidious Poison in the Body Politick?*, 1 J. PUB. ECON. 199 (1972). Some scholars have proposed alternatives. See, e.g., G. Munda, P. Nijkamp, and P. Rietveld, *Information Precision and Multicriteria Evaluation Methods*, in EFFICIENCY IN THE PUBLIC SECTOR: THE THEORY AND PRACTICE OF COST-BENEFIT ANALYSIS 43 (Alan Williams and Emilio Giardinia eds. 1993); V. Kerry Smith, *A Conceptual Overview of the Foundations of Benefit-Cost Analysis*, in BENEFITS ASSESSMENT: THE STATE OF THE ART 13, 27-31 (Judith D. Bentkover, Vincent T. Covello, and Jeryl Mumpower eds. 1986). We discuss these and other alternatives in Part III.

apparently the case that when people are lost, they tend to walk in circles. Whatever the true path out of the woods, walking in circles is inferior to walking in a straight line. So a method that enabled a person to avoid walking in circles (for example, walk toward some landmark) is superior to no method at all. Similarly, CBA may enable agencies to avoid certain errors—like the tendency to exaggerate certain benefits and to ignore certain costs—without actually telling the agency whether a project is desirable. The critics of CBA, however, reply that while this is possible, it is not likely. For them, CBA is like a method for leaving the woods that is no better than random. For example, a method that tells one to rely on a divining rod, or one that tells one to walk in figure eights, is no better than random. The argument that CBA is better than an alternative cannot be made independently of a theoretical defense of it and a comparison to its rivals. That is the burden of the Part III.

III. RECONCEPTUALIZING COST-BENEFIT ANALYSIS

In this Part, we argue that CBA is properly conceptualized as a *welfarist decision-procedure*. We defend the following claims: (1) the effect of a governmental project on overall well-being is a morally relevant, if not morally decisive feature of the project; (2) CBA is plausibly the decision-procedure best justified in light of overall well-being, in a significant fraction of agency choice situations. We do not claim that CBA is *always* best justified in light of overall well-being. On the other hand, CBA has certain distinct advantages, relative to other decision-procedures, which suggest that it will be routinely if not universally appropriate.

Our conception of CBA has a number of salient features. First, we conceive of CBA as a *decision-procedure*, not a criterion of moral rightness or goodness. The fact that a project has a positive sum of compensating variations says nothing at all, even *prima facie*, about the moral worth of the project. The criterion of overall well-being, and other true moral criteria, are conceptually distinct from the sum-of-CVs test. For example, as we discussed above, the project winners might be rich and the project losers might be poor, such that the winners would be willing to pay large sums in dollars for trivial welfare benefits, and the losers would require smaller sums for larger welfare harms. Yet it is a large mistake to leap from the premise that CBA lacks

bedrock moral status, to the conclusion that agencies ought not employ CBA as a decision-procedure. CBA might be *sufficiently* accurate in tracking the welfare effect of projects that—notwithstanding the conceptual slippage between CBA and overall well-being—it is the best procedure for agencies to use, given the relative cheapness and transparency of CBA.

Second, our conception severs any link between CBA, and the two purported moral criteria most familiar to modern economists: namely, Kaldor-Hicks and Pareto-superiority. Again, the considerations motivating this stance were discussed in Part II. Kaldor-Hicks is not, in truth, a moral criterion; the fact that the winners from the project could compensate the losers does not, without more, mean that the project is a good one, even *prima facie*. As for the criterion of Pareto-superiority: although that *does* mark out something morally significant, it is a criterion of limited scope. It leaves unranked projects that have both welfare winners and losers, as agency projects typically do.

Rather, our conception ties CBA to a much older criterion than Kaldor-Hicks or Pareto-superiority, a criterion with an impressive philosophical pedigree: *overall well-being*. Modern economists are often uncomfortable with that criterion, because of the standard claim about the impossibility of interpersonal welfare comparisons. That claim is wrong, as we shall argue at length below. All of us regularly compare welfare gains and losses, across persons, and indeed a conception of welfare that precluded such comparisons would be unreasonable.

We should emphasize that by asserting the possibility of interpersonal welfare comparisons, and the moral relevance of overall well-being, we are *not* committing ourselves to the truth of utilitarianism. Utilitarianism is the view that overall well-being is morally decisive: the only important feature of a project is its effect on aggregate welfare. Our view, a considerably weaker one, is that overall well-being is morally *relevant*: government should choose a welfare-improving project, but all things considered, non-welfarist considerations (for example, distributive or deontological considerations) may properly lead to the ultimate rejection of that project. CBA is a decision-procedure by which agencies implement one of the several normative criteria that, together, determine the all-things-considered normative status of the project.

Finally, as should become clearer below, our conception of CBA

is distinctive because we detach CBA from *preferences*. The concept of preference is foundational to modern economics, both positive economics (where actors are assumed to maximize satisfaction of their preferences) and normative or welfare economics (where satisfying preferences is assumed to make persons better off). This may or may be not appropriate in the first case B we take no position on that, because our project here is normative, not positive B but we do claim that the preference-satisfaction view of welfare is wrong. More precisely, as we shall explain, the standard, economic view that makes welfare depend upon the satisfaction of unrestricted preferences is quite misguided. But this is not a difficulty for CBA, because CBA is agnostic across different conceptions of well-being. The idea of measuring a project's effect on overall well-being by monetizing the effect on each individual, and then aggregating, does not presuppose an unrestricted preference-based view of well-being, or indeed any preference-based view at all. To put the point another way: the "compensating variations" summed to determine the overall costs and benefits of a project should be defined, not as a person's willingness to pay or accept (which presupposes a preference-based view of welfare), but as her *welfare equivalent* (leaving open what the right theory of welfare is).⁶¹

We proceed as follows. Section A describes different theories of well-being, and criticizes the unrestricted preference-based view. Section B and C address, respectively, two different objections to the purported moral criterion of overall well-being. The first is the conceptual objection that interpersonal welfare comparisons are impossible. The second is the normative objection that, conceding the possibility of such comparisons, overall well-being is nonetheless morally irrelevant. In Section D, we flesh out the distinction between decision-procedures and moral criteria, and discuss how CBA (understood as a welfarist decision-procedure) should be defined. In Section E, we compare CBA with alternative decision-procedures B for example, with direct implementation of the welfare criterion; with nonaggregative procedures; with unidimensional procedures; and with

⁶¹ Actually, this is just a first cut at our position. More precisely, CVs should be defined as either welfare equivalents, or as WTP/WTA to the extent the latter measure tracks welfare equivalents with sufficient accuracy and is cheaper, more transparent, more reliably implemented, and so forth. See *infra* text accompanying note ____.

other multidimensional procedures B and highlight CBA's advantages, in accuracy, cost, and transparency, relative to these other procedures, in light of overall well-being. Finally, in Section F, we briefly discuss the problem of bringing nonwelfarist criteria, such as deontological or distributive criteria, to bear on agency choices.

Before we proceed, a terminological point is in order. The term "welfarist" is used, throughout this Part, as a synonym for "relative to overall well-being." "Welfarism," in our terminology, is the view that overall well-being is morally relevant. Hence our description of CBA as a welfarist decision procedure. The terms are sometimes, not always, used this way in the philosophical literature; in any event, this is the way they will be used here.

A. Well-Being: Desire, Pleasure and Objective Value

Philosophers standardly divide theories of well-being into three types: desire-based theories, objective-list theories, and hedonic theories.⁶² Theories of well-being can be used for various purposes. We are interested, in this Article, in their use for illuminating the problem of *comparative* well-being.⁶³ CBA is a technique by which agencies compare two or more options: options such as issuing various possible rules, or issuing no rule; spending money in various possible ways, or not spending it; and so on. As we explained in Part II, cost-benefit analysts refer to an agency's options as either "projects" or the "status quo," and we will follow that locution here. What must be true of the project world, for a given person A, such that A is better off there than in the status quo? Desire-based theories, objective-list theories, and

⁶² For overviews of the philosophical literature on well-being, see JAMES GRIFFIN, *WELL-BEING* 7-72 (1986); L.W. SUMNER, *WELFARE, HAPPINESS, AND ETHICS* 45-137 (1996). Sometimes the trichotomy of welfare theories is drawn differently from the way we have just drawn it: as a trichotomy of desire-based theories, objective-list theories, and *mental-state* theories, where the last category in turn is given two subcategories: hedonic theories, and mental state theories (such as Sidgwick's) that define welfare in terms of desirable mental states rather than positive feeling-tones. See SUMNER, *supra*, at 90-91 (discussing Sidgwick's view). Since we are trying to draw a sharp distinction between welfare theories that rest upon desires, and theories that do not, we include the Sidgwickian variant within our category of desire-based theories.

⁶³ On the importance of comparisons for justified choice, see generally Ruth Chang, *Comparison and the Justification of Choice* 146 U. PA. L. REV. 1569 (1998).

hedonic theories give different answers to this question. A desire-based theory says that it is a *necessary* condition for A to be better off with the project, that A desire the project over the status quo. By contrast, both objective-list and hedonic theories deny that A's desiring the project world is a necessary condition for her comparative benefit. Instead, for the objective-list theorist, A's welfare in the project world, as compared to the status quo world, depends upon the balance of objective goods that A realizes in the two worlds.⁶⁴ And for the hedonic theorist, A's comparative welfare depends upon the balance of pleasurable mental states that A realizes in the two worlds.

What is a "desire"? By that term, we simply mean what philosophers call a "pro-attitude": that is, some kind of *propositional attitude* with a *favorable* valence.⁶⁵ "Desires," thus defined, are a *generic* category which include such specific pro-attitudes as wants, hopes, wishes, favorable judgments, preferences, lusts, likes, and so forth. What these all have in common is that they take states-of-affairs as their objects. In this important way, desires are different from non-propositional mental states such as physical pleasures and pains. I feel thirsty, or hot, or itchy. I desire *that* I learn this musical composition, or that the parade take place, or that we go on a ski trip. Further, desires are favorable rather than unfavorable. I like learning musical compositions; I detest reading novels. Both liking and detesting are propositional attitudes, but only liking is a desire because only liking is favorable.

⁶⁴ What these goods are is a matter for further debate, within the family of objective-list theories, but standardly they are taken to include such goods as knowledge, personal relationships, play, the experience of beauty, the accomplishment of worthy goals, and physical fitness. For some specific lists of objective values, see JOHN FINNIS, *NATURAL LAW AND NATURAL RIGHTS* 85-90 (1980); GRIFFIN, *supra* note __, at 67-68; GEORGE SHER, *BEYOND NEUTRALITY: PERFECTIONISM AND POLITICS* 199-201 (1997).

⁶⁵ On propositional attitudes, see JAEGWON KIM, *PHILOSOPHY OF MIND* 13-14 (1996); on the "valencing" of propositional attitudes, see RICHARD BRANDT, *A THEORY OF THE GOOD AND THE RIGHT* 24-45 (1998 ed.) By contrast with Brandt, however, we think that valence need not be defined in terms of choice. For example, I might retrospectively endorse some state of affairs involving me, even though I didn't choose it, and even though I endorse it just because it happened to me serendipitously rather than through my choice. Cf. SUMNER, *supra* note __, at 122-37 (discussing retrospective endorsement).

We should emphasize that our use of the term *desire*, here, is stipulative. The term *desire*, in ordinary English, has certain connotations (for example, affective connotations) that we do not intend. Our definition of *desire* tracks the broader usage common in the philosophical literature. Any pro-attitude is, for our purposes, a *desire*. And any theory of well-being that makes the satisfaction of some pro-attitude a necessary condition for A's welfare is, by our definition, a "desire-based" theory.

The predominant theory of well-being within welfare economics is a specific variant of a desire-based theory. This specific variant makes well-being depend upon *preferences*, which are one type of pro-attitude.⁶⁶ How are preferences different from judgments, endorsements, wishes, and other pro-attitudes? Although economists differ over this issue, in the standard textbook treatment of CBA, a person has a preference for some good if he chooses that good rather than alternatives. Thus, preferences, unlike certain other pro-attitudes, are conceptually connected to choice: if I prefer P over S, then I am necessarily disposed to choose P over S.⁶⁷ Moreover, preferences, unlike certain other pro-attitudes, have no necessary emotional or affective

⁶⁶ On the nature of preference, see, e.g., SUSAN HURLEY, *NATURAL REASONS* 55-83 (1989); Richard J. Arneson, *Liberalism, Distributive Subjectivism, and Equal Opportunity for Welfare*, 19 PHIL. & PUB. AFF. 158, 161-64 (1990); Arthur Ripstein, *Preference*, in VALUE, WELFARE AND MORALITY 93-111 (R.G. Frey & Christopher Morris eds., 1993). For a full analysis of the related concept of a "want," see ROBERT AUDI, *ACTION, INTENTION, AND REASON* 35-55 (1993). Critical discussions of the link, within welfare economics, between preference and well-being are provided by SUMNER, *supra* note __, at 113-22; DANIEL HAUSMAN & MICHAEL S. MCPHERSON, *ECONOMIC ANALYSIS AND MORAL PHILOSOPHY* 71-83 (1996); and Tyler Cowen, *The Scope and Limits of Preference Sovereignty*, 9 ECON. & PHIL. 253 (1993). We do not necessarily endorse all the criticisms voiced by these authors, but instead focus specifically on the point that well-being is not equivalent to the satisfaction of *unrestricted* preferences. See *infra* text accompanying note __.

⁶⁷ See, e.g., HAUSMAN & MCPHERSON, *supra* note __, at 28 ("We regard Q's preference ranking as a subjective state of Q that, along with Q's beliefs, explains her choices"). This is not quite the same as saying that Q's preferences are "revealed" in her choices, since, for example, Q's beliefs may be mistaken. See *id.* (criticizing revealed preference view). For our purposes, we need not decide whether the conceptual connection between preference and choice is properly captured by a "revealed preference" view, or in some other way.

component: I can prefer P over S without having a strong feeling about either option. Further, preferences, unlike certain other pro-attitudes, have no necessary cognitive component: I can prefer P over S without reason. Finally, and most significant, the standard economic theory makes well-being depend upon *unrestricted* preferences. The theory says that, assuming constant preferences, A is better off with the project if, all things considered, A prefers the project world (or would do so given sufficient information).

The standard economic theory is wrong, and we reject it. It is wrong because A might prefer the project to the status quo for all manner of reasons, including but not limited to his welfare.⁶⁸ For example, A might prefer the project to the status quo because he believes that the project is morally required, even though A also believes he would be personally better off with the status quo. Imagine a project that redistributes resources from the rich, including A, to the needy poor. To insist that the project here improves A's own welfare is confused and mistaken, for several reasons: first, considerations of morality and welfare notoriously can conflict, and yet the unrestricted preference-based theory says that necessarily considerations of morality and welfare do not conflict where A's preferences track what morality requires;⁶⁹ second, what underlies preference-based and more broadly desire-based theories is the accurate intuition that A is, in some way,

⁶⁸ We are hardly the first to articulate the point that the satisfaction of unrestricted preferences should not be conflated with welfare. Others who have made the same point include: SUMNER, *supra*, note __, at 134-35; Allan Gibbard, *Interpersonal Comparisons: Preference, Good and the Intrinsic Reward of a Life*, in FOUNDATIONS OF SOCIAL CHOICE THEORY 165, 173-75 (Jon Elster & Aanund Hylland eds., 1986); John Broome, *Choice and Value in Economics*, 30 OXFORD ECON. PAPERS 313 (1978) []; Mark Carl Overvold, *Self-Interest and the Concept of Self-Sacrifice*, 10 CAN. J. PHIL. 105 (1980); Amartya Sen, *Rational Fools: A Critique of the Behavioral Foundations of Economic Theory*, 6 PHIL. & PUB. AFF. 317 (1977); David Sobel, *On the Subjectivity of Welfare*, 107 ETHICS 501 (1997).

⁶⁹ As Overvold puts it: "[S]uppose we accept the prevailing account of self-interest or personal welfare. Then we will have to say that any act that is voluntary and informed is thereby in the agent's self-interest. But self-sacrifice requires that the act be voluntary, informed, and contrary to the agent's self-interest. Thus accepting the prevailing account of self-interest makes the concept of self-sacrifice incoherent by making it logically impossible that there are every genuine instances of self-sacrifice." Mark Carl Overvold, *Morality, Self-Interest, and Reasons for Being Moral*, 44 PHIL. & PHENOMENOLOGICAL RESEARCH 493, 499 (1984).

sovereign with respect to his own well-being, and yet here the welfare economist insists that A is better off with the project even though A's own judgment is that, for himself, the status quo is better.

But there *is* an important insight underlying the standard economic theory, and that is, as we have just said, that A is in some way sovereign with respect to his own well-being. Objective-list and hedonic theories ignore this crucial point. Both say, in different ways, that A can be better off in the project world even if his various pro-attitudes all point in favor of the status quo world. Start with a hedonic theory. A hedonic theory identifies one or more kinds of "pleasures"—technically, one or more types of non-propositional mental states characterized by positive feeling-tones—such that if A realizes the better mix of pleasures in the project world, he is better off there. As Charles Sumner explains, in describing hedonism:

There is a core of physical pleasures which are the counterparts in every respect of physical pains: they have a purely organic basis, they are often localized in one part of the body, they can have a quite specific duration, they vary in intensity, and we employ a similar vocabulary for describing the way they feel. The paradigm instances are the pleasures caused by stimuli such as scratching an itch, being massaged, taking a hot bath, quenching a thirst, using a recreational drug, urinating, defecating, and sexual arousal and orgasm. What these sensations have in common, in virtue of which we distinguish them from physical pain, is just the fact that they feel good. When asked to characterize the peculiar feeling tone of sensory pleasure (or pain) we find, like Bentham, that we have little to say.⁷⁰

But it is a mistake to think that "pleasure," without desire—that is, without a favorable attitude on the part of the person who experiences the positive feeling tone—suffices to make that person better off. Persons can and do judge that pleasure, beyond a certain point, would be excessive, and purely from the point of view of self-interest. I can decide that the third glass of wine, or the fourth sweet, or whatever, is

⁷⁰ SUMNER, *supra* note __, at 106. See also KIM, *supra* note __, at 13 (distinguishing between "sensations" and "propositional attitudes").

just too much for me; I favor having only two glasses, or three sweets, and if so I am at least no worse off with that.⁷¹

A similar objection can be leveled against objective-list theories. To see the point clearly, we must distinguish between objective goods that entail pro-attitudes, and objective goods that do not. "Recreation" or "play" is an objective value that, presumably, entails a pro-attitude. I'm not truly playing a game if I'd prefer not to be. A theory of well-being predicated solely on these kinds of desire-entailing goods is a desire-based theory, by our scheme. Rather, our objection is to the kind of theory that is not thus predicated. Take goods such as "knowledge" or "musical accomplishment." You can know a lot without wanting, wishing, hoping, or preferring to know a lot; you can be an accomplished musician (say, a prodigy forced into music by an overbearing parent) but prefer, want, wish, judge, and like a life without music. An objective-list theorist, by our categorization, says that it is possible for A to be benefitted by the project, even though all his desires point in favor of the status quo, if, further, he realizes the right mix of (non-desire-entailing) objective goods in the project world. This kind of theory, like hedonism, fails to respect A's own point of view about what makes him better off. The prodigy who realizes great musical feats, but truly would prefer watching sit-coms, is not comparatively better off for his musical accomplishments. The world may be, but he is not. At best he is neither better off, nor worse off, as between the accomplishment world and the sit-com world.

In short, the right theory of well-being is (some variant) of a

⁷¹ The sophisticated hedonist might try to solve these counterexamples by specifying mixes of pleasures so as to take account of considerations of balance, excessiveness, and so forth: 3 parts gustatory pleasure and 3 parts sexual arousal are better than 10 parts gustatory pleasure and 2 parts sexual arousal. Yet this more refined hedonic theory is still open to the objection that it fails to respect A's point of view. Imagine that A better realizes the specified mix of pleasures in the project world, but he nonetheless prefers, judges, endorses, and otherwise desires the status quo. (This is a conceptual possibility because, again, "pleasures" are non-propositional. For the pleasures that go into a hedonic theory, it is no entailment of A's experiencing such pleasures that he, further, desire them.) In this sort of case, we propose, A is not better off with the project. Perhaps he would be better off if the project also brought with it an adaptive desire, a desire for the project-induced mix of pleasures; but if it does not, then the project does not improve A's welfare. At most he is neither comparatively better off, nor comparatively worse off, as between the project and status quo. See Parfit ___.

restricted desire-based theory. It is a necessary condition, for A to be benefitted by a project, that he actually desire the project at some point in time. A further, necessary, condition is that the desire concern A's own life: this is what we mean by "restricted." The mere fact that some desire of A's has been satisfied, by the project, is insufficient to ensure that his welfare is improved, even *prima facie* B since A might desire the project on moral, or other disinterested grounds. To be sure, how to provide a more precise and persuasive account of this "restriction" remains a large and unsolved problem within the philosophical literature on well-being.⁷² (Clearly, to say that A's desires are properly "restricted" if they concern A's own welfare, and that it is both necessary and sufficient to improve A's welfare that his thus-restricted desires be satisfied, would be circular.) We need not and will not attempt to solve that problem here. Suffice it to say that we find persuasive the propositions that (a) A cannot be benefitted by a project if he never desires it, or comes to desire it; and (b) simply satisfying any desire of A's cannot be enough to benefit him. A "restricted, desire-based theory" of well-being is simply a theory that takes both propositions to be true.

It also bears emphasis how large the category of such theories is. The concept of "desire," as we mean it here, is generic. A desire is simply a propositional attitude with a favorable valence. Restricted desire-based theorists can disagree with each other about what kind of pro-attitudes (preferences versus judgments versus *ex post* endorsements) are relevant. They can disagree about how to handle the notorious problem of changing desires: Can A be benefitted by the project if he desires it before it occurs, but not afterwards? If he desires it at some moments in his life, or only if he desires it at a majority of moments? Finally, a restricted desire-based theory says that satisfying a restricted desire is a *necessary* condition for a project to improve someone's well-being. It need not be sufficient. A desire-based theorist can insist that, for a project to comparatively benefit A relative to the

⁷² Mark Overvold has made a sustained attempt to provide a non-circular account of the restriction. See Overvold, *supra* note __; Overvold, *supra* note __; Mark Carl Overvold, *Self-Interest and Getting What You Want*, in THE LIMITS OF UTILITARIANISM (Harlan Miller and William Williams ed. 1982). It is far from clear whether he succeeds. See Sobel, *supra* note __, at 795 n.24.

status quo, A must desire the project (at the right time) *and* one or more the following must hold true.

— *The Experience Requirement.* Gardening is my sole avocation. I work long and hard to till an exotic garden, designed to foster exotic and sensitive plants that will take many years to bloom, if they ever do. I strongly desire that the plants eventually bloom and endure. Indeed they do, but by that point I have moved far away, and never learn of the garden's success. My desire has been satisfied (without my experiencing that.) Has my well-being improved? Maybe not.⁷³

— *Informed Desires.* Looking back over my professional life, I contemplate an intense and exciting year I spent in a government-funded research lab. I say to myself, "That year was fabulous. We solved tough problems; we worked together as a team; I felt great." Unbeknownst to me, the results of our research were funneled to, and instrumental in the success of, a secret weapons program. Has my well-being been improved by the research year, notwithstanding my uninformed approval? Maybe not.⁷⁴

— *The Affect Requirement.* I work long and hard to complete a project. At the end of the project's completion, I feel nothing—nothing at all. It's not that, in general, I lack the capacity to feel. My affective capacity is of the ordinary type. Rather, for whatever reason, this particular project leaves me empty. Has its completion improved my well-being? Maybe not.⁷⁵

— *Objective Value.* I'm obsessed with romance films. I make sure to see every new release within this category, and spend

⁷³ See SUMNER, *supra* note __, at 128 (discussing importance of experience).

⁷⁴ See Sobel, *supra* note __, at 792 n.15 (noting that "[a] truly impressive and diverse list of contemporary ethicists have found [a] full-information [desire-based] account of well-being congenial" and citing Brandt, Hare, Griffin, Rawls, Gauthier, Darwall, and Harsanyi).

⁷⁵ See SUMNER, *supra* note __, at 138-56 (discussing importance of affect).

most of my free time watching videos of famous and not-so-famous romance films from the past. After great effort, I have succeeded in memorizing the names of the actors and characters, and the plot names, of every romance film ever made. And I enjoy all this very much. Would I be even better off if I had spent all of this avocational energy on high art, or philosophy, with equal enjoyment? Perhaps so.

Here, too, we intend to remain agnostic about specifics. Whatever the specific desire-based theory, it will remain true that overall well-being is both non-empty and morally important, and that CBA is a plausible decision-procedure by which to implement the criterion of overall well-being.

B. The Possibility of Interpersonal Welfare Comparisons

A common objection to desire-based theories is that they depend on the possibility of interpersonal comparisons of utility while being inconsistent with any reasonable theory of interpersonal comparisons of utility. Indeed, as we discussed in Part II, economists developed the compensation tests because they (i) rejected Bentham's hedonic utilitarianism in favor of preference-based utilitarianism, but (ii) did not believe that interpersonal comparisons of utility were possible or manageable under a theory of preference-based utilitarianism. Because we reject the Kaldor-Hicks justification of CBA, while endorsing a desire-based theory of welfare, we must explain why we do not think that the problem of interpersonal comparisons of utility is insurmountable.

Part of what animates the traditional view is a general skepticism about the truth content of moral and evaluative statements. Such skepticism is, now, a distinctly minority position within analytic philosophy⁷⁶ and—more to the point—incompatible with the claim that, truly, satisfying preferences or desires improves welfare. A second and more philosophically robust basis for the economist's traditional concern about interpersonal comparisons—a concern quite distinct

⁷⁶ See STEPHEN DARWALL, ALLAN GIBBARD & PETER RAILTON, *Toward Fin de Siecle Ethics: Some Trends*, in MORAL DISCOURSE AND PRACTICE: SOME PHILOSOPHICAL APPROACHES 3 (1997) (providing overview of contemporary metaethics).

from general skepticism about moral or evaluative truth—is the worry that preferences or desires as such are simply *ordinal rankings*.⁷⁷ Where A ranks the project over the status quo, and B ranks the status quo over the project, and we know nothing more, we have no basis for saying that the project improves or degrades overall well-being.

However, desire-based theories of well-being are capable of incorporating quite a bit more than information about ordinal rankings. In the interpersonal context, desire-based theories *must* incorporate more welfare-relevant information about A and B than the bare fact of their ordinal rankings. In day-to-day life, we routinely make judgments of overall well-being, comparing losses to some of our friends, colleagues, or family members with benefits to others. It is a condition on the validity of a welfare theory that it warrant some judgments of this kind. As Daniel Hausman puts it: “[I]f a conception of well-being does not permit one to make interpersonal comparisons in an acceptable way, than that conception is itself unacceptable.”⁷⁸ The fact that a theory of well-being that (a) makes well-being consist in the satisfaction of bare desires or preferences, and that (b) incorporates no mechanism for translating the ordinal rankings constituted by desire or preference into some interpersonally comparable form, leads to (c) the impossibility of interpersonal comparisons, does *not* imply that such comparisons are indeed impossible. Rather, it implies that the theory is wrong!

So how should comparisons of well-being be made, on a beefed-up desire-based theory? We are asking, here, not how real-world agencies should in fact pick out welfare-improving projects, but rather how the concept of an interpersonal comparison should be constructed—that is, how an epistemically perfect agency with no procedural costs would do so. Initially, a project will produce Winners and Losers. Winners are those who are better off with the project, as compared to the status quo. Losers are those who are worse off with

⁷⁷ See, e.g., Daniel Hausman, *The Impossibility of Interpersonal Utility Comparisons* 104 MIND 475-77 (1995).

⁷⁸ *Id.* at 473. For “[i]nterpersonal comparison are an ineliminable part of human life.” *Id.* at 489. See also John C. Harsanyi, *Morality and the Theory of Rational Behavior*, in UTILITARIANISM AND BEYOND 39, 49 (Amartya & Bernard Williams eds. 1982) (“In everyday life we make, or at least attempt to make, interpersonal utility comparisons all the time”).

the project, as compared to the status quo. A Winner must, at a minimum, desire the project at some point in his life; a Loser must, at a minimum, desire the status quo. That is just what a desire-based theory means. How, more specifically, to classify persons as Winners and Losers depends on the specifics of the theory and such problems as desires that change over time.⁷⁹

In any event, the agency must compare welfare gains to the Winners with welfare losses to the Losers. The conceptually thorny problem, of course, is how precisely to do *that*.⁸⁰ Broadly speaking, there seem to be two views within the literature. One view is that the interpersonal comparison can and should hinge on the degree to which the Winners' and Losers' welfare-constitutive desires are satisfied or frustrated, *independent of the extent to which Winners and Losers improve or decline with respect to criteria of objective value or hedonic tone*.⁸¹ The idea, in economists' terms, is to *cardinalize* A's and B's preferences so that A's gain from the project can be compared with B's loss. But how to cardinalize? Perhaps the most famous answer is that given by Harsanyi, with his construct of extended preferences. Harsanyi, essentially, proposes this: (1) for every outcome O_k and every person in the population A_i , permute the persons and outcomes; (2) imagine yourself to be an impartial spectator, comparing person-state O_k-A_i with person-state O_1-A_i (that is, "being in state O_k with A_i 's preferences" as against "being in state O_1 with A_i 's preferences"); (3) from this impartial point of view, create a preference ranking, the so-called extended-preference ranking, for all person-states, which (Harsanyi proposes) should be the same for everyone, and which should

⁷⁹ More precisely, the agency would need to sort all persons into Winners, Losers, and Neutrals, where Neutrals are those who are neither better off nor worse off—which in turn means either that they are precisely as well off in both world-states, or that they are incomparably well off as between the states. Neutrals would then be ignored at the second stage of the interpersonal comparison. The possibility of Neutrals is a technical issue that need not be further discussed here, because that possibility does not bear on the key problem of interpersonal comparisons, namely how to compare Winners' gains to Losers' losses.

⁸⁰ See generally Peter J. Hammond, *Interpersonal Comparisons of Utility: Why and How They Are and Should be Made*, in INTERPERSONAL COMPARISONS OF WELL-BEING 200 (Jon Elster and John Roemer eds. 1991); Hausman, *supra* note

⁸¹ For an overview of such constructs, see Hausman, *supra* note ____.

respect A_i 's ordinary preferences in comparing two person-states that both involve her; (4) determine the extended preference ranking for the O_k lottery (that is, for a $1/n$ chance of O_k permuted with each of the n persons in the population) as against the O_1 lottery; 5) use this extended lottery preference as the social welfare function for comparing O_k to O_1 .⁸²

The Harsanyi construct remains quite controversial. It is unclear whether that construct truly solves the problem of rendering individual rankings of outcomes comparable, across persons, independent of an appeal to the objective value or hedonic tone of the outcomes. Harsanyi proposes that we each develop impartial or extended preferences over person-states, and that these extended rankings converge on a single extended ranking. But why believe that the rankings will converge—why *will* we all rank O_k - A_i over O_1 - A_j —absent some independent standard by reference to which person-states can be ranked? To quote one critic of Harsanyi: “[T]he relevant [extended] preference must be purged of the judge’s own personal tastes, attitudes, feelings, moral views and so on. The problem is, then, how, after that sort of purging [and absent further reference to objective criteria] can form any sort of preference at all.”⁸³

⁸² See Harsanyi, *supra* note __; see also John A. Weymark, *A Reconsideration of the Harsanyi-Sen Debate on Utilitarianism*, in INTERPERSONAL COMPARISONS OF WELL-BEING, *supra* note __, at 255-320; Kenneth Arrow, *Extended Sympathy and the Possibility of Social Choice*, 7 PHILOSOPHIA 223 (1978). A less elaborate construct than Harsanyi's is the so-called "zero-one" rule, first suggested by Isbell. For each person, construct the familiar *intrapersonal* cardinal index of welfare known as the von Neumann-Morgenstern index (that is, the index that assigns numbers to outcomes such that a person's preferences over these outcomes are tracked by the index numbers, *and* a person's preferences over lotteries of the outcomes are tracked by the expected value of the index numbers for the outcomes). Make sure, additionally, that for each person the index number 1 is assigned to her highest outcome, and 0 is assigned to the lowest. Then, Isbell proposes, the interpersonal comparison of two outcomes, O_i and O_j , is quite simple: simply determine the numerical difference between the two outcomes on each person's zero-one index, and aggregate. This proposal is discussed in Hausman, *supra* note __; and Hammond, *supra* note __, at 215-16. See also Ken Binmore, GAME THEORY AND THE SOCIAL CONTRACT: PLAYING FAIR 282-96 (1994).

⁸³ James Griffin, *Against the Taste Model*, in INTERPERSONAL COMPARISONS OF WELL-BEING, *supra* note __, at 53. For a similar criticism, see Hausman, *supra* note __, at 477-78. Similarly, Isbell's zero-one rule can be criticized because it

It is beyond the scope of this Article to address the controversy over the Harsanyi construct, and over other attempts to compare the degree to which outcomes satisfy different persons' desires, independent of criteria of objective value or hedonic tone. We need take no position on the controversy here. For even if it turned out to be true that a purely desire-based construct for interpersonal comparisons were impossible, that would not entail the impossibility of interpersonal comparisons as such. It would simply entail that interpersonal comparisons do depend, in part, on criteria of objective value or hedonic tone.

James Griffin is perhaps the leading exponent of this latter view, within the family of desire-based theorists. As Griffin explains:

[Consider] Mill's interpersonal comparison of Socrates and the Fool. The Fool attaches no value to Socrates' life. Socrates attaches no value to the Fool's life. How would each decide how relatively well off they are? . . . What Socrates [or the Fool] needs to make is a judgment of a very different sort from what we ordinarily understand by a personal preference. . . . Socrates [or the Fool] should need to know, primarily, what made life valuable. He should have to appeal to his understanding of what humans, or sometimes humans of a certain type, are capable of, and of the various peaks that human life can reach. Then he should have to decide how close he and the Fool came to some peak. *What he should not particularly need to consult is the phenomenological "feel" of their experience, nor their personal tastes and attitudes, nor his own preferences about landing in the one sort of life or the other.*⁸⁴

We emphasize that Griffin *is* a desire-based theorist, indeed one of the

relies upon information about how persons intrapersonally rank outcomes under risk—each person's index number for O_k , or for O_l , depends upon her comparative ranking of that outcome, relative to various lotteries of other outcomes—while what we want to know interpersonally is how the occurrence of O_k *with certainty* compares to the occurrence of O_l *with certainty*.

⁸⁴ GRIFFIN, *supra* note __, at 117 (emphasis added). For a similar construct, albeit within the context of a more objectivist view of well-being, see Thomas M. Scanlon, *The Moral Basis of Interpersonal Comparisons*, in INTERPERSONAL COMPARISONS OF WELL-BEING, *supra* note __, at 39-44.

leading ones.⁸⁵ But a desire-based position is consistent with Griffin's further claim that it is objective values, and not strength-of-desire, that furnishes the basis for interpersonal comparison.

To return to our schema of project Winners and Losers: it is a necessary condition for a person to be a project Winner that, at some time, she desire the project over the status quo. The basic sorting of persons into these categories depends, in part, on their desires. *This* is where the desire-based theorist and other theorists of well-being disagree. If you never desire the project, you can't be a Winner; you're a Loser (or a Neutral), and if the last you're ignored entirely in aggregating welfare losses and gains.⁸⁶ But desiring the project is simply a *necessary* condition, within a desire-based theory, for a person to be a Winner; it need not be a sufficient condition, as we have already discussed in the intrapersonal context. And, relatedly, in the interpersonal context, the desire-based theorist can say that the comparative welfare of the Winners and Losers depends upon the objective values or hedonic tones that they, comparatively, realize.

We are not endorsing Griffin's view here. Rather, we cite Griffin's work to help make evident the point that the following two propositions are consistent: (1) any reasonable theory of well-being must make A's desiring the project a necessary condition for A's benefit; and (2) any reasonable theory of well-being must warrant interpersonal comparisons, i.e., statements to the effect that welfare gains to the winners outweigh or are outweighed by welfare losses to the losers. Desire-based theories are reasonable, in both senses. They, definitionally, satisfy the first proposition. And they can warrant interpersonal comparisons B perhaps in the manner that Harsanyi proposes, but at a minimum in the manner that Griffin does.

C. The Moral Relevance of Overall Well-Being

Our defense of cost-benefit analysis rests on the premise that agencies should, within certain constraints, promote the overall well-being of citizens. Not everyone, however, agrees that overall well-being is an appropriate moral criterion.

How can one deny the moral relevance of overall well-being? First, one can do so by holding the kind of moral view that Shelly

⁸⁵ GRIFFIN, *supra* note __, at 7-72. []

⁸⁶ On "Neutrals," see *supra* note __.

Kagan aptly calls “minimalist.”⁸⁷ The “minimalist” believes that persons have no moral reason whatsoever to make the world better (at least not a reason of the kind that can ground a moral requirement,⁸⁸ and that can thereby warrant *government* in taxing and coercing persons merely for the sake of world-improvements). Libertarianism is the most famous, modern variant of minimalism.⁸⁹ Libertarianism says that persons *are* morally obliged to comply with certain “deontological” constraints (the constraints against killing, stealing, assaulting, defrauding, and so on, as exemplified by the prohibitions of the criminal law), but insists that (a) these constraints cannot be reinterpreted in the form of “consequentialist,” i.e., world-improving, requirements; and that (b) persons are under no moral requirement beyond the requirement of compliance with deontological constraints.⁹⁰

Minimalism, of the libertarian variant or any other, is arguably at odds with commonsense moral views.⁹¹ Even more clearly, minimalism is inconsistent with the scope of modern government. Consider, for example, an environmental agency. The libertarian holds, in effect, that the agency can only proscribe those acts endangering the environment that (assuming *mens rea* on the polluter) would justify criminal punishment quite apart from the agency’s regulation. But surely there are large parts of the clean air, clean water, and endangered species laws that proscribe actions not independently punishable under the criminal law—actions that do not, apart from the agency’s regulation, fall within some traditionally criminal *actus reus*. Or consider a food and drug agency. It violates neither the deontological constraint on fraud, nor any other deontological constraint, to sell to

⁸⁷ See SHELLY KAGAN, *THE LIMITS OF MORALITY* 1-10 (1989).

⁸⁸ See *id.* at 64-70 (distinguishing between moral reasons and moral requirements.)

⁸⁹ See ROBERT NOZICK, *ANARCHY, STATE AND UTOPIA* (1974).

⁹⁰ The constraint against killing, as proposed by the libertarian, is not a consequentialist requirement, because it violates the constraint even if that violation serves to prevent more killings. For an accessible overview of the consequentialism/deontology distinction, see SHELLY KAGAN, *NORMATIVE ETHICS* 25-105 (1998). For a more technical discussion, see, e.g., David McNaughton & Piers Rawling, *Agent-Relativity and the Doing-Happening Distinction*, 63 *PHIL. STUDIES* 167 (1991).

⁹¹ See KAGAN, *supra note* ___, at 16.

consumers a fully and accurately labeled pharmaceutical product, complete with warnings, that (as it happens) consumers are likely to misuse. And yet the licensing of drugs for prescription use only is a central function of the FDA. Finally, consider the antitrust laws. An agreement among firms to set the price of a product at a particular level is, for the minimalist, just another free contract. And yet it is a core tenet of modern antitrust law that price-fixing is illegal.⁹²

What about a non-minimalist critique of CBA? The non-minimalist critic agrees that consequentialist criteria B criteria that mark out different ways in which the world goes better or worse B *do* bear upon an agency's choice of projects, but argues that "making the world better" (Ahaving good consequences") and "promoting overall well-being" are wholly distinct.⁹³ Within the modern philosophical literature, there are two ways that philosophers have drawn this distinction—two ways in which consequentialism has been severed from the criterion of overall well-being. The first way is to argue for some consequentialist standard of distribution other than aggregation, for example, for equalizing well-being, or for maximizing the well-being of the person with the lowest welfare. This is a line of argument that goes back to Rawls, who famously asserts in *A Theory of Justice* that social contractors behind the veil of ignorance would choose a maximin standard for distributing primary goods; and more recently has been developed by egalitarian theorists such as Dworkin,⁹⁴ Cohen,⁹⁵ Temkin⁹⁶ and Van Parijs.⁹⁷ We will call this, a bit roughly, the

⁹² See Nozick — (on monopoly). Our argument here is, of course, only the briefest sketch of the case against minimalism. We will not try to build a fuller case because minimalism has not figured significantly within the scholarly debate about CBA; the critics of CBA have, at least implicitly, agreed that consequentialist criteria of some kind do bear upon agency choices.

⁹³ The distinction between consequentialism and utilitarianism has been drawn very clearly in the recent philosophical literature. See, e.g., SAMUEL SCHEFFLER, *THE REJECTION OF CONSEQUENTIALISM*, 1-40 (rev ed. 1994). The position that overall well-being lacks even moral relevance, let alone moral decisiveness, is the limiting point of nonutilitarian consequentialism.

⁹⁴ See Ronald Dworkin, *What is Equality? Part 2: Equality of Resources*, 10 PHIL. & PUB. AFF. 283 (1981).

⁹⁵ [cite from recent P & PA].

⁹⁶ See LARRY TEMKIN, *INEQUALITY* (1993).

⁹⁷ See PHILIPPE VAN PARIJS, *REAL FREEDOM FOR ALL: WHAT (IF ANYTHING)*

egalitarian criticism of overall well-being (and CBA). The second way to sever consequentialism and “overall well-being” is to bracket the distributive issue, but argue that well-being as such is morally unimportant and that some more basic element or prerequisite for life-prospects, for example, the satisfaction of needs, is what really counts morally. This second line of argument goes back, again, to Rawls—with his claim that the principles of justice govern the distribution of *primary goods*, that is, resources for welfare, and not welfare itself—and more recently has been developed in different ways by Dworkin,⁹⁸ Scanlon,⁹⁹ Nagel,¹⁰⁰ Cohen¹⁰¹ and Sen.¹⁰² We will call this second line, a bit roughly, the *resourcist* criticism of overall well-being (and CBA). For example, Nagel argues:

If you and a stranger have both been injured, you have one dose of painkiller, and his pain is much more severe than yours, you should give him the painkiller—not for any complicated reasons, but simply because of the relative severity of the two pains, which provides a neutral reason to prefer the relief of the more severe. The same may be said of other basic elements of human good and will.

But many values are not like this. Though some human interests (and not only pleasure and pain) give rise to impersonal values, I now want to argue that not all of them do. If I have a bad headache, anyone has a reason to want it to stop. But if I badly want to climb to the top of Mount Kilimanjaro, not everyone has a reason to want me to succeed. I have a reason to try to get to the top, and it may be much

CAN JUSTIFY CAPITALISM? (1995).

⁹⁸ See Dworkin, *supra* note ____.

⁹⁹ See Scanlon, *supra* note ____.

¹⁰⁰ See Thomas Nagel, *Autonomy and Deontology*, in CONSEQUENTIALISM AND ITS CRITICS 142 (Samuel Scheffler ed. 1988).

¹⁰¹ See G.A. Cohen, *Equality of What? On Welfare, Goods and Capabilities* in THE QUALITY OF LIFE 9 (Martha Nussbaum & Amartya Sen eds. 1993).

¹⁰² See, e.g., AMARTYA SEN, *INEQUALITY REEXAMINED* (1992); Amartya Sen, *Capability and Well-Being*, in THE QUALITY OF LIFE 9 (Martha Nussbaum & Amartya Sen eds. 1993).

stronger than my reason for wanting a headache to go away, but other people have very little reason, if any, to care whether I climb the mountain or not.¹⁰³

Since the climber's ascent of Kilimanjaro would, in turn, improve the climber's well-being—he desires the ascent with sufficient information, it is objectively good, as a kind of accomplishment, and so on—and Nagel would seemingly concede as much, what Nagel is claiming here is that only certain prerequisites or elements of well-being, such as physical pleasures and pains, constitute a good feature of world-states that persons in general (and the government) have a reason to promote.¹⁰⁴

How should the welfarist defender of CBA respond to the egalitarian and the resourcist? Let us start with the egalitarian. The egalitarian (as we use that description here) claims that the aggregate amount of well-being, or the aggregate amount of some prerequisite or element of well-being, e.g., physical pain or pleasure, is morally irrelevant. She identifies some set of non-aggregative consequentialist criteria, such that these criteria, along with deontological criteria, *exhaust* the moral considerations bearing upon governmental choices. By “non-aggregative,” we mean that the criteria do not take the following form: they do not mark out an increase in well-being, or its prerequisites or elements, as a moral improvement independent of further conditions, in particular, independent of the level of welfare or wealth of the persons who gain and lose. Consider, for example, the following types of egalitarian goals, regularly discussed in the literature: equalizing welfare, maximizing the welfare of the least-well-off (maximin), and bringing persons up to a minimum level of welfare. If a project increases A's welfare more than it reduces B's, and does nothing else, then the equalizer counts the project as morally bad if B is poorer than A, and as morally good if (roughly) A is poorer than B. The maximin theorist will count the project a matter of moral indifference, unless A or B or perhaps both are members of the least-

¹⁰³ Nagel, *supra* note __, at 145-46.

¹⁰⁴ To put the point another way: on the resourcist view, the fact that a project increases overall well-being, without increasing (or otherwise properly distributing) the specified prerequisites or elements of well-being, gives government no reason whatsoever to approve the project.

well-off group. Finally, the minimum-welfare theorist will count the project a matter of moral indifference, unless it brings A above the minimum level or forces B below that level.

Egalitarianism, in the robust sense delineated here, is inconsistent with the structure and processes of government as we know it. The welfarist will happily concede that some agencies, such as welfare agencies, should be solely concerned with implementing criteria of fair distribution; and that other agencies, such as environmental, or antitrust, or food-and-drug agencies, might take account of non-aggregative criteria along with considerations of overall well-being. By contrast, the egalitarian needs to argue that *every* agency should pursue the equalization of welfare, or the maximization of the welfare of the least-well-off, or some other such non-aggregative goal, as its *sole* intrinsic aim. In the case of an environmental agency, for example, this would mean that the decision to prohibit an environmental pollutant imposing health risks upon a particular segment of the population, should depend crucially and solely on the level of welfare (or wealth) of that segment of the population, as compared to the level of welfare (or wealth) of the consumers, workers, and shareholders who would bear the costs of controlling the pollutant. In the case of an antitrust agency, egalitarianism would mean that the decision to bar price-fixing in a particular industry (thereby redistributing welfare from the firms' shareholders to its consumers) should depend crucially and solely on the level of welfare (or wealth) of the shareholders versus that of the consumers. Clearly, the decision-procedures that EPA, the FDA, the FTC, and other ordinary agencies employ, look nothing like this. Such agencies routinely pursue goals—be these goals such as minimizing pollution simpliciter, or minimizing pollution in a cost-justified way, or reducing safety risks to a de minimus level, or minimizing the total number of deaths, or maximizing economic surplus (the area between supply and demand curves) in an industry—that have an aggregative structure.

The egalitarian might respond by saying that ordinary agencies pursue aggregative goals only in the service of deeper, non-aggregative ones. Maximizing collective welfare works to the advantage of the least-well-off, or to persons below a minimum level of welfare, and conversely is morally justifiable only insofar as it does so. This strikes us as an implausibly limited account of the moral justification behind ordinary agencies. Imagine that the shareholders of a firm that emits

pollution, or sells risky products, are middle class. The pollution, or products, cause pain and injury to breathers or to consumers who happen to be even richer than the firm's shareholders. The pollution or the product defect can be eliminated at a very small cost to the firm and, ultimately, the shareholder. Should it be eliminated? The egalitarian is committed to saying no, unless the pain or injury to the rich breathers or consumers is sufficiently large to bring their level of lifetime welfare below that of the shareholders. Or to build upon Nagel's example: if you have one dose of a painkiller, which can reduce some slight discomfort of yours, and the person next to you is suffering intense, traumatic pain, you should give him the painkiller even if your future income is \$30,000/year and his is \$200,000.

In short, aggregative considerations are surely *relevant* to judgments of good consequences. Saying that is consistent with the claim that non-aggregative, specifically egalitarian goals are also relevant,¹⁰⁵ indeed that some non-aggregative goals take lexical priority over welfarist or resourcist goals. But even if it is true that certain non-aggregative goals are morally relevant or even take lexical priority over the welfarist criterion, CBA is plausibly *one part* of the total decision-procedure that government properly employs. For example, as between two projects, neither of which affects the extent of poverty or violates deontological constraints, "overall well-being" plausibly provides the decisive moral consideration between them, and if CBA tracks overall well-being then the agency ought to employ that decision-procedure in choosing between the two projects.

Now for the debate between resourcist and welfarist. The resourcist concedes that certain aggregative goals have intrinsic importance for government. "Agencies *do* maximize," the resourcist will allow. "But what they properly maximize is some *element or prerequisite* of well-being, rather than well-being itself." Agencies maximize the satisfaction of needs, or (on Nagel's view) the relief from physical pain; as between two world-states that equally satisfy egalitarian criteria and deontological criteria, the resourcist will choose that state where the aggregate amount of the relevant element or prerequisite is higher. The difficulty with this view is explaining why, among the elements or prerequisites of B's well-being that make a moral claim on A, and that

¹⁰⁵ See, e.g., KAGAN, *supra* note __, at 25-70.

justify a governmental project harming A and helping B, only certain elements or prerequisites of welfare do so. For presumably those elements and prerequisites derive their moral force on A from the fact that, in turn, they are elements and prerequisites of B's *welfare*. So why not simply say that B's welfare makes a moral claim on A? As James Griffin nicely puts the point:

A group of scholars may, with full understanding, prefer an extension of their library to exercise equipment for their health. And part of what makes us think that basic needs, such as health, are more closely linked to obligation than are desires is that basic needs seem the "bread" of life and desires mere "jam." But an extension to the scholars' library may not seem like "jam" to them. On the contrary, if the scholars preference is sufficiently informed then the library is of greater value to them. But then to maintain that needs create obligations where mere desires do not, or that they create stronger obligations, is to say that we have an obligation, or a stronger one, to the scholars to give them what they themselves value less, which would be odd.¹⁰⁶

A similar welfarist critique could be made of Nagel's distinction between pains and pleasures and "mere" desire-satisfaction. If producing pain causes you a welfare setback of X, and painlessly frustrating your desires causes you a welfare setback of X+K, then why should the first setback make a claim on me, but the second not all?

One possible resourcist response is to advert back to welfare levels—for example, to define a "need" as any setback to a person that puts her beneath a certain minimum level of welfare, rather than (as in Griffin's example) a health or safety setback to a person, the scholar, who already is well enough off. "Maximizing the satisfaction of needs" then means minimizing the extent to which people are below that basic level. But this strategy turns resourcism into a type of egalitarianism; it smuggles a non-aggregative structure into the resourcist goal, through the definition of a "need." We have already argued that egalitarianism does not plausibly exhaust the set of goals

¹⁰⁶ GRIFFIN, *supra* note __, at 45.

that governmental agencies properly pursue; some of their proper intrinsic goals are aggregative, and the question here is whether resourcism can provide a persuasive picture of *those* goals. We suggest not. If “maximizing the satisfaction of needs” or “maximizing the relief from pain” is intrinsically important—important beyond the egalitarian benefits that happen to flow from agency pursuit of these goals—then the importance thereof derives from the fact that unsatisfied needs and unrelieved pains inhibit welfare. But so, too, do (some) unfulfilled desires.

D. Decision Procedures and Moral Criteria

Having established that the promotion of a desire-based conception of overall well-being is an appropriate goal of agencies, our next task is to link this goal with cost-benefit analysis. Our claim is that CBA is an appropriate means, or decision-procedure, for achieving the morally desirable goal of promoting overall well-being.

The distinction between a *criterion of moral rightness or goodness*, such as the criterion of “overall well-being,” and a *decision-procedure justified in light of that criterion*, is an intuitive and (as far as we know) uncontested distinction. A criterion of moral rightness or goodness marks out the properties of some action, some state of affairs, or some other thing that constitute, or partly constitute, the moral status of that thing. For example, a consequentialist criterion identifies some feature of a state-of-affairs such that a state-of-affairs possessing that feature, or possessing it to a greater extent, is better or worse (at least holding other things equal) than a state-of-affairs lacking that feature, or possessing it to a lesser extent. “Overall well-being” is, we have argued, such a criterion. A state-of-affairs with a higher level of aggregate well-being is, everything else equal, better than a state-of-affairs with a lower level of aggregate well-being. Derivatively, an action (specifically, a governmental project) leading to a state-of-affairs with a higher level of aggregate well-being is, *ceteris paribus*, better than an action leading to a state-of-affairs with a lower level of aggregate well-being.

By contrast, a decision-procedure justified for an agent in light of some (consequentialist) moral criterion is the following: the procedure for choosing between actions such that the agent’s use of that procedure leads to the best consequences, as measured by that criterion. Specifically, the decision-procedure justified for a governmental agency, in light of the criterion of overall well-being, is the procedure for

choice between projects such that the agency's use of that procedure maximizes overall well-being. Clearly, this concept of the welfare-justified decision-procedure is distinct from the concept of overall well-being. Whether a project improves or degrades well-being, relative to the status quo, is one thing; what steps the agency should take, in deciding between the project and the status quo, is quite another. There is no contradiction in saying that (1) the best project for an agency to choose is the project that maximizes well-being; and (2) the best way for an agency to *decide* which project to choose is *not* to attempt to identify the project that maximizes overall well-being.¹⁰⁷

What drives this conceptual wedge between a criterion of moral rightness or goodness, specifically the criterion of overall-being, and a decision-procedure justified in light of that criterion? There are a number of factors that do so. First is the possibility of *epistemic imperfection* on the part of the agent. The agent might make mistakes in deciding what the criterion requires; and it might further be the case that some other (morally irrelevant) standard is both reasonably well-correlated with the criterion, and less subject to mistaken application by the agent, such that the best decision-procedure for the agent is *not* direct implementation, but rather the implementation of the correlated standard.

Another factor that helps drive the wedge between moral criteria and decision procedures, is the factor of *cost*. Assume that the agent is epistemically perfect. Given sufficient time and effort, she will always pick out the project that meets the applicable criterion. Even so, the process of doing so might consume lots of time and effort B resources that, employed in other ways, might improve the satisfaction of the criterion even more than their use by the agent for making this choice.

Yet another factor concerns the *faithfulness* of the agent and, relatedly, the opacity of the procedure B the ease with which third parties can verify that the agent has faithfully attempted to follow, or succeeded in following, the procedure. Assume that the agent is epistemically perfect and, further, that direct implementation is cheap. Nonetheless, the agent might be unfaithful. If instructed that she is under a legal and sanction-backed obligation to directly implement the

¹⁰⁷ The distinction between decision-procedures and moral criteria is discussed in DAVID BRINK, MORAL REALISM AND THE FOUNDATIONS OF ETHICS 216-17 (1989).

criterion at stake, the agent might aim at other goals (for example, self-regarding goals) and yet credibly claim (given the opacity of direct implementation) that she has complied with the instruction. Even if an alternative decision-procedure is more expensive than direct implementation, and the agent is less epistemically reliable in following the alternative procedure, it still might be the case that B given the transparency of the alternative procedure B instructing the agent to follow the alternative procedure has better results overall, in terms of the underlying criterion, than instructing the agent to engage in direct implementation.

The distinction we are drawing here is familiar to legal scholars. The distinction is parallel to the distinction, in the legal literature, between standards and rules. “Standards” are the moral criteria bearing on some actor’s choice. But the best way to implement a given standard might be to instruct the actor to apply a rule that tracks the standard well enough, given the cheapness with which the actor can apply the rule, the ease with which his compliance can be monitored, and so on. A “rule” is a simply another term for what we are calling a decision-procedure: some specification of actions, states and contexts that the actor should actually follow in making his choice.

The distinction between criteria of moral rightness and goodness, and morally justified decision-procedures, has important implications for the legitimacy of CBA. It implies that the legitimacy of CBA is a moral problem *and* an institutional one, not a moral problem alone. Critics of CBA have often assumed that, by undermining the moral status of CBA B by showing why a positive sum-of-CV’s marks out nothing of bedrock moral importance about a project, even *prima facie* B they have succeeded in making their case against the procedure. Thus the dominant focus, in the critical literature, on the features of CBA that no respectable moral criterion would (allegedly) possess: on Scitovsky reversals; on the sensitivity of CVs to the winner’s or loser’s endowment; on the moral irrelevance of wealth-maximization *per se*.¹⁰⁸ But it is a mistake to leap from the existence of these features, to

¹⁰⁸ See, e.g., BOADWAY AND BRUCE, *supra* note __, at 262-72; JULES COLEMAN, *Efficiency, Utility and Wealth Maximization*, in *MARKETS, MORALS AND THE LAW* 95 (1988); HAUSMAN & MCPHERSON, *supra* note __, at 93-99; Charles Blackorby & David Donaldson, *The Case Against the Use of the Sum of Compensating Variations in Cost-Benefit Analysis*, 23 *CAN. J. ECON.* 471 (1990);

the conclusion that it is wrong for CBA to be institutionalized as the method by which agencies choose between projects. Notwithstanding the intermittent occurrence of reversals, or of cases where (by virtue of CBA's sensitivity to endowments) a welfare-degrading project has a positive sum-of-CVs, CBA might be *sufficiently* accurate in tracking overall well-being and *sufficiently* cheap and transparent, that it turns out to be the decision-procedure best justified in light of overall well-being. Or to put the point reciprocally: given the moral criterion of overall well-being, it is a further and institutional question whether the welfare-maximizing decision-procedure, for agencies, is: (1) the procedure of *direct implementation* (where agencies do indeed attempt to identify welfare-maximizing projects); (2) CBA, or some refinement thereof; (3) some nonaggregative procedure, for example looking to technical feasibility or social norms; (4) some unidimensional, aggregative procedure, for example so-called "risk-risk" balancing; or (5) some multi-dimensional, aggregative procedure other than CBA, for example QUALY-based assessment.

In Section E, below, we will undertake a comparison of these different types of procedures, and will argue that CBA (or some refinement thereof) possesses certain advantages, such that it is plausibly the welfare-maximizing procedure for agencies to employ in a significant portion of their choice situations. Before we undertake this comparison, however, it is important to broaden the definition of CBA. Once we understand that CBA is properly conceptualized as a decision-procedure, not a basic moral criterion, there is no reason to insist that CBA is strictly equivalent to the traditional sum-of-compensating variations test. Rather, we suggest, CBA should be understood as a *family* of money-based decision procedures, including but not limited to the strict sum-of-CVs test. We think it implausible that CBA, strictly defined, is in fact a welfare-justified decision procedure. (Among other things, the cost of individualizing CVs would be overwhelming. In practice, as we noted in the Introduction, agencies do not actually determine, for each project, what each affected person's CV for that project is; rather, agencies use an average value for the affected population, or rely on statistical techniques to estimate the

David Copp, *The Justice and Rationale of Cost-Benefit Analysis*, 23 THEORY AND DECISION 65 (1987); Ronald Dworkin, *Is Wealth a Value?*, 9 J. LEGAL STUD. 191 (1980).

range of variation of CVs across the population.) Our claim is that one or another *refinement* to the strict sum-of-CVs test is plausibly welfare justified for a significant fraction of agency choice situations. And this claim fits actual practice, given the reliance of agencies upon refinements, not the strict test. In the remainder of this Section, we discuss the different ways in which traditional CBA might be refined, to suit its possible role as a welfarist decision-procedure.

One important kind of refinement concerns the definition of a CV. Compensating variations (CVs) are traditionally equated with WTP or WTA. Person A's CV for a project, it is traditionally stipulated, is the amount that he would be willing to pay (if a winner) or willing to accept (if a loser) such that, if paying or paid that amount in the project world, A would neither prefer the project to the status quo nor vice versa. But this definition of compensating variation ties CBA to an incorrect theory of well-being: an unrestricted preference-based theory. Imagine that A slightly prefers the project, for himself, but also judges it to be morally wrong and further judges the project world to be morally worse than the status quo regardless of how much money *he* is paid there. (The project, let us imagine, eliminates an endangered species, or a wilderness area, so as to build a road that will make A's daily routine slightly more convenient, but which he finds morally objectionable because of its environmental impact.) Then A's genuine CV for the project would seem to be a positive number, but his CV, as traditionally defined in terms of WTP/WTA, is negative infinity!¹⁰⁹ Or imagine a case in which A strongly prefers the status quo for himself, but the project is a paternalist project (for example, banning narcotics) that is welfare-justified precisely because persons like A mistake their own well-being. Then A's CV, as traditionally defined in terms of WTP/WTA, will be a large negative number B some narcotics users might demand large payments in the project world

¹⁰⁹ Indeed, this phenomenon—the deviation between a person's traditionally defined CV and her true welfare equivalent for a project that has a small welfare effect on her, but to which she has a strong moral objection—seems to arise frequently in agency practice. Respondents often react to contingent-valuation surveys in ways that seemingly reflect their moral views, for example by providing a very large or small number, or by simply refusing to answer. See ROBERT CAMERON MITCHELL & RICHARD CARSON, USING SURVEYS TO VALUE PUBLIC GOODS: THE CONTINGENT VALUATION METHOD 30-38 (1989).

before they count it as equivalent to the status quo B but A's genuine CV would seem to be either zero or a positive number. The narcotics-eliminating project improves A's welfare (if he eventually comes to prefer not using narcotics) or at least does not change it (if his desires never change); it does not harm A, notwithstanding his mistaken belief that the status quo is better.

The appropriate redefinition of CBA is straightforward. The "compensating variation" seeks to capture, in dollar terms, the effect of the project on each person's well-being. It seeks to measure, on a dollar scale, the change in A's welfare produced by the project. It should thus be defined, not as A's willingness to pay or accept, but as her *welfare equivalent*. A's welfare equivalent is the amount of money that, paid to or from A in the project world, makes her precisely as well off there as in the project world. While the concept of "willingness to pay" and "accept" is committed to a particular, and incorrect theory of well-being, the concept of welfare-equivalent is agnostic across theories of well-being and would seem to be the right foundation for CBA whatever the correct theory of well-being turns out to be. We have argued, specifically, that some variant of a restricted, desire-based theory is the correct theory. So, for example, a project that satisfies A's unrestricted desires, but either (a) does not satisfy or frustrate his restricted desires, or (b) does so, but does not fulfill further appropriate conditions (e.g., an experience requirement, or a value requirement) necessary for a welfare effect upon A, has a welfare equivalent of zero. More generally (leaving aside, for the moment, certain technical problems such as the problem of uncompensable losses), A's welfare equivalent will be positive if the project has a positive welfare effect on A; negative if the welfare effect is negative; and zero if the welfare effect is nil. This will be true whatever specific variant of a restricted, desire-based theory turns out to be correct. (As we have already explained, we are not committed here to a particular variant.)¹¹⁰

¹¹⁰ As we briefly mentioned in the text above and in the introduction to this Part, the concept of welfare-equivalent is an appropriate foundation for CBA even if the correct theory of well-being turns out to be a hedonic or objective-list theory. CBA is no less coherent, on these theories, than on a desire-based theory, although of course the specific amount of A's welfare equivalent will be different. On a hedonic theory, that amount is the amount that, paid to or from A in the project world, makes his level of hedonic tone there just the same as in the status

Now, the defender of willingness-to-pay and -accept has a sophisticated response to the proposal that CBA should be defined in terms of welfare equivalents. The response runs as follows: “The right theory of well-being is complicated and controversial. Although A’s welfare-equivalent and her WTP/WTA might indeed diverge, we should not instruct agencies to determine welfare equivalents. Faithful agencies are likely to make frequent mistakes in determining what A’s welfare equivalent is, insofar as it diverges from her WTP/WTA; and the instruction to aggregate welfare-equivalents rather than WTP/WTA will give unfaithful agencies an increased opportunity for shirking, given the relative opacity of the concept of >welfare-equivalent.’ Thus, notwithstanding the failure of an unrestricted preference-based theory of well-being, the welfare-maximizing decision procedure for agencies to follow (even on the correct, restricted theory) turns out to be the procedure of aggregating WTP/WTA, or some other procedure grounded on WTP/WTA, not a procedure grounded on welfare equivalents. For WTP/WTA is *sufficiently* accurate in tracking welfare equivalents, and is also cheaper to apply, more transparent, etc.” This sophisticated response is absolutely correct. Once we reconceptualize CBA as a welfarist decision-procedure, it is incorrect to assume that the definition of CVs will directly incorporate the right theory of well-being. CVs might be defined in a different way (e.g., as WTP/WTA) that tracks the correct theory sufficiently

quo; on an objective-list theory, it is the amount that, paid to or from A in the project world, makes his overall realization of objective welfarist values just the same as in the status quo world. However, because we argue in this article for a desire-based view of well-being, we do not further pursue this point.

The point that A’s WTP/WTA is distinct from his welfare equivalent has been made by Peter Railton, in *Benefit-Cost Analysis as a Source of Information About Welfare*, in VALUING HEALTH RISKS, COSTS, AND BENEFITS FOR ENVIRONMENTAL DECISION MAKING 55, 71-72 (P. Brett Hammond & Rob Coppock eds., 1990). Cf. Richard H. Pildes and Cass R. Sunstein, *Reinventing the Regulatory State*, 62 U. CHI. L. REV. 1, 82 (1995) (“[C]ontingent valuation methods still suffer from the public/private valuation distinction. What people should be willing to pay to eliminate certain conditions for themselves, and how they think public resources should be allocated, remain distinct questions. . . . [W]hat matters are lay valuations about public choices, not those about self-regarding, private choices”). We disagree: in the context of determining overall welfare, what matters just are the “self-regarding,” i.e., welfare-relevant impacts of the project on each person.

well. So both WTP/WTA and welfare-equivalents are possible variants of the concept of CV; which variant is best, in light of overall well-being, is an institutional and empirical question that cannot be determined in advance. On the other hand B and this is our ultimate point here—the traditional equation of CV with WTP/WTA is too narrow. CBA might, depending on the facts, need to be (globally or locally) redefined as the sum of welfare-equivalents rather than the sum of WTP/WTA. This is one possible, and significant dimension of refinement of traditional CBA.

A second dimension of possible refinement to CBA concerns the problem of *undefined CVs*. This is a large and important issue, one that the traditional literature on CBA has wholly overlooked, and that we do not have the space to consider here in detail, but surely bears mention.¹¹¹ Economists traditionally assume that compensating variations are precise and unique. There will be a unique amount of money that, paid to or from A in the project world, precisely counterbalances the project's welfare effect on him; further, larger amounts will make him strictly better off or worse off, and smaller amounts strictly worse or better off. Or so it is assumed. But behind the assumption of precision lies the contestable premise that all world-states can be ranked as better, worse, or precisely equal with respect to A's well-being. This premise is contestable, insofar as some (indeed, an increasing number of) philosophers, economists, and legal scholars believe that world-states can be *incomparable* with respect to well-being and other moral criteria, i.e., neither better, nor worse, nor precisely equal.¹¹² (If incomparability can obtain, then there may well be projects such that, for no dollar payment to or from A along with the project is A precisely as well off in both worlds, but for some dollar payment A is incomparably well off. In such a case, there will exist no precise welfare equivalent, for A, for the project, but only a rough welfare substitute.) And even if full welfare-comparability of world states holds true, it does not follow that A's welfare equivalent will be unique. The highly contestable premise behind uniqueness B behind the assumption that one, and only one, dollar amount counterbalances the project's welfare effect on A B is the premise that the correlation

¹¹¹ The issue is treated at greater length in Matthew Adler, *Incommensurability and Cost-Benefit Analysis*, 146 U. PA. L. REV. 1371, 1391-98 (1998).

¹¹² See *id.* at 1401-08 (discussing incomparability).

between dollars and well-being is continuous, monotonically increasing, and unbounded. Counterexamples to this premise include: (a) uncompensable losses, i.e., welfare differences between the project and status quo that are too large for any dollar payment in the project world to repair; (b) lumpy compensation, e.g., for the loss of nonmarket goods, where a range of dollar payments in the project world would leave A precisely as well off there as in the status quo world.

The refinements to CBA, responsive to the problems just mentioned, are refinements that change in one way or another the concept of CV, so that A's "CV" will remain well-defined notwithstanding incomparability, discontinuity, non-monotonicity, etc. Rather than being defined as the unique amount of money that precisely counterbalances A's welfare loss or gain, the CV might be redefined, for example, as follows: (a) in the case where A is a loser, the smallest amount of money sufficient either to compensate precisely A, or to overcompensate A, or to leave her incomparably well off, or failing that the lower limit of the amounts that precisely compensate A, overcompensate A, or leave A incomparably well off; (b) in the case where A is a winner, the largest amount of money that still leaves A better off, or precisely as well off, or incomparably well off, or failing that the upper limit of the dollar amounts that do so. As between several refinements to CBA, all of which succeed in preserving well-defined CVs notwithstanding incomparability, etc., the best refinement is of course the one that, used by agencies, will maximize welfare.

A third dimension of refinement to CBA, already mentioned, concerns *the degree of individualization of CVs*. However CVs are defined—whether as simple WTP/WTAs, or simple welfare equivalents, or with some amendment to the basic idea of WTP/WTA or welfare-equivalent designed to deal with the problem of undefined CVs B it is implausible that agencies should literally determine, for each person, what her CV is, and then aggregate. Some method for approximating the sum-of-CVs (e.g., determining average CV and then estimating the variation of CV across the population) will surely be warranted, if CBA is warranted at all.

Finally, CBA might perhaps be refined to correct its *endowment-dependence*. to compensate for the declining marginal welfare productivity of dollars or, equivalently, for the fact that richer persons

tend to have larger monetary CVs for a given change in interpersonal welfare than poorer persons. The standard suggestion, here, is to weight CVs by a factor inversely proportional to the wealth of the person affected, and then sum weighted CVs.¹¹³ Welfare economists have not yet, in fact, been successful in producing a practicable weighting factor; but it would be premature to insist that CBA will never be successfully refined along this dimension.¹¹⁴

* * *

In sum, CBA, traditionally defined as the sum of CVs (WTP/WTA) should be refined in at least four ways: by redefining CVs as welfare equivalents; by redefining them to correct for the possible absence of unique and precise welfare equivalents or WTP/WTA; by reducing individualization; and by correcting for endowment-dependence. To be sure, the possibility of such refinements does not, yet, show that refined CBA is indeed a welfare-maximizing decision procedure. To do that, we need to compare CBA, and its refinements, with alternative decision-procedures, including both direct implementation and others. That is the task to which we now turn.

E. Evaluating CBA

In this Section, we compare direct implementation of the welfare criterion with various decision procedures actually employed by, or proposed for, regulatory agencies: procedures such as CBA (or some refinement), risk-risk balancing, feasibility-based assessment, norm-based assessment, and others. We shall distinguish between *non-aggregative* and *aggregative* procedures; in the latter category, between *unidimensional* aggregative and *multidimensional* aggregative procedures; and, in the last category, between the three leading candidates, namely CBA (or a refinement), *QUALY-based assessment*,

¹¹³ See, e.g., Copp, *supra* note __, at 77-79.

¹¹⁴ If no refinement along this dimension eventuates, the upshot is *not* that CBA must be abandoned, but rather that it must be confined to choice situations where endowment-dependence does not cause too great a degree of inaccuracy. B specifically, situations where the wealth distribution within the group of project Winners does not differ too much from the distribution within the group of Losers.

and the kind of procedure proposed, among others, by Thomas Scanlon, Cass Sunstein, and Richard Pildes, which we shall call *direct multidimensional assessment*.¹¹⁵ Our claims are as follows. First, direct implementation of the welfare criterion is not a viable decision procedure; that procedure is hugely expensive and highly opaque and unreliable, particularly given the amount of individualized welfare information required by any reasonable construct for making interpersonal comparisons. Second, multidimensional aggregative procedures, including CBA, will generally be more accurate than nonaggregative procedures and unidimensional procedures, although they will also be more expensive to implement and, to some extent, more susceptible to agency error and more opaque. If agencies can be suitably monitored, (some kind of) multidimensional aggregative procedure will typically be appropriate for evaluating projects that appear to have large welfare impacts. Finally, CBA (or some refinement) offers distinct advantages, relative to QUALY-based assessment and direct multidimensional assessment, with respect to accuracy, transparency, or error rate, depending on the choice situation.

Our claims in this Section are couched in vague terms like “significant,” and are qualified by tempering adjectives such as “plausibly” and “likely.” Such imprecision and qualification are inevitable, given the informal and context-independent cast of the analysis attempted here. The analysis is necessarily informal, given the large number of specific decision procedures being compared. (A formal analysis would need to specify which refinement of CBA was being considered, as against which version of direct multidimensional assessment, say, and in light of which particular construct for making interpersonal comparison.) And it seems appropriate to us to undertake this kind of comparison in an overarching way, without reference to a particular regulatory context B such as the context of environmental law, or antitrust law, or workplace risk regulation B given the inchoate state of the debate about CBA. If a scholarly consensus emerges that CBA is properly conceptualized as a welfarist decision-procedure, more specific and formal work will then need to be done. Our goal in this Article is the logically prior one of showing that CBA is properly thus

¹¹⁵ More precisely, these are the procedures standardly employed or proposed as background procedures, for use by agencies when statutory and other legal requirements become indeterminate.

conceptualized B which means showing, negatively, that CBA is not a moral criterion and, affirmatively, that overall well-being *is* a moral criterion such that (some refinement of) CBA is plausibly justified as a decision-procedure in light of that, in a significant fraction of agency choice situations.

Why not directly implement the criterion of overall well-being? By *direct implementation*, we mean the procedure where an agency expends resources up to the point of zero marginal epistemic gain (up to the limits of human knowledge) in order to determine how the project and the status quo compare with respect to the best construct for making interpersonal comparisons. That is, the agency is instructed to make all efforts that will increase its degree of justified belief in the comparative worth of the project and status quo, with respect to overall well-being. No proxy is assigned the agency, in lieu of the best interpersonal construct. And no constraints are placed on the agency's gathering of information, other than the epistemic constraint that the information must have some relevance with respect to that construct.¹¹⁶ Clearly, direct implementation would be hugely expensive and highly opaque (at least for agency projects that affect more than a few individuals), and for these reasons is not a viable welfarist decision-procedure. This is true whether the best interpersonal construct is a desire-based construct (such as Harsanyi's construct), or a construct based on objective values or hedonic tones (such as that suggested by James Griffin).

Consider Harsanyi's proposal, which is the leading candidate for a desire-based interpersonal construct. (By "desire-based," here, we mean a construct that makes the comparison of welfare across persons depend upon the extent to which winners and losers satisfy their welfare-constitutive desires, independent of their improvement or deterioration with respect to criteria of objective value or hedonic tone).

¹¹⁶ This stipulation is crucial. We are trying to draw a clean line between direct implementation, and procedures that sacrifice accuracy (more precisely, the accuracy that would be achieved by a perfectly reliable and faithful agent) for the sake of cost, transparency, reliability and so on. The cleanest way to do that is to define "direct implementation" as the procedure where the agent tolerates *no* sacrifice in epistemic gain for the sake of cost or other values. The agent takes whatever steps are warranted, by her lights, to produce the most accurate measurement of overall well-being.

In the case of a comparison of a single project O^* and the status quo O , recall that Harsanyi's proposal amounts to this: (1) for each outcome, O^* and O , and every person in the population A_i , permute the persons and outcomes; (2) imagine yourself to be an impartial spectator, comparing person-state O_k-A_i with person-state O_l-A_j (that is, "being in state O_i with A_i 's preferences" as against "being in state O_j with A_j 's preferences"); (3) from this impartial point of view, create a preference ranking, the so-called extended-preference ranking, for all person-states, which should be the same for everyone, and which should respect A_i 's ordinary preferences in comparing A_i-O and A_i-O^* ; (4) determine the extended preference ranking for the O^* lottery (that is, for a $1/n$ chance of O^* permuted with each of the n persons in the population) as against the O lottery; 5) use this extended lottery preference as the basis for comparing O^* to O . Directly implementing the Harsanyi construct would involve a huge amount of individualized information. For each person A_i , the agency would need to know how she ranks the project, relative to the status quo. Further, the agency would need to know everything else about the project's effect on the person that would bear upon an impartial ranking of $O-A_i$ as against O^*-A_i , and of O^*-A_i as against $O-A_i$. And quite apart from these informational demands, Harsanyi's construct would give agencies significant scope for shirking and error, since there is often no clear answer to how impartial spectators would rank person-states (assuming, as Harsanyi does, that there is at least a right answer to that question). Not surprisingly, the Harsanyi construct has *only* been discussed in the economic and philosophical literature as a plausible analysis of the concept of "overall well-being."¹¹⁷ We know of no one who suggests that regulatory agencies should actually apply it.

Similar difficulties of implementation beset reasonable constructs for interpersonal comparison that depend, as does James Griffin's, on objective or hedonic criteria and ignore "desire-satisfaction" apart from that. First, any such construct will need to be (at bottom) multidimensional. There are multiple objective goods that go to human

¹¹⁷ Discussions in the philosophical literature include: HURLEY, *supra* note __, at 103-11; Griffin, *supra* note __, at 52-56; Hausman, *supra* note __, at 477-78; Scanlon, *supra* note __, at 22-38. Discussions in the economic literature include Binmore, *supra* note __.

welfare, and multiple kinds of pains and pleasures.¹¹⁸ Call each dimension of objective value or hedonic tone, D_i . Second, any such construct will need to allow that the effect of a project, on a person, with respect to a particular dimension of value, is individualized. A beautification project, for example, does not make every person's life more beautiful to the same degree; the impact on each person depends upon her natural appreciation for beauty, her training, her temporal and spatial proximity to the project, and so on. To put the point formally: the project's effect with respect to D_i cannot be represented as $D_{i,project}^k$ minus $D_{i,status\ quo}^k$ but must instead depend upon the aggregate of $D_{i,project}^k$ minus $D_{i,status\ quo}^k$, where k represents a particular person and where these numbers can and do vary between persons. Further, the *tradeoff* between the different dimensions of objective value or hedonic tone will itself be personalized. Imagine a project that increases beauty in the world for A by 1 unit, but consumes resources that would otherwise fund 2 units of scientific accomplishment by A. If A is an artist, the project might still be welfare-improving; if A is a scientist, it probably will not be.¹¹⁹ Finally, the proper assessment of how each person fares, both with respect to each D_i and all-things-considered, would be highly indeterminate. Here, as with the Harsanyi exercise of developing extended preferences over person-states, there would be ample room for agency mistakes, and for an agency's pursuit of its own projects under the cover of making "reasonable" judgments of value or hedonic tone.

In short, direct implementation is not a live option, for agencies. That alone is an important result. The standard criticism of CBA, for its characteristic inaccuracy B for permitting small welfare gains to the rich to outweigh large welfare gains to the poor B implicitly compares CBA to the perfectly accurate procedure, namely direct

¹¹⁸ See sources cited *supra* note __ (providing and defending lists of objective values). Cf. SEN, *supra* note __, at 23-26 (recognizing multidimensional cast of the assessment of equality); THOMAS HURKA, PERFECTIONISM 84-98 (1993) (recognizing multidimensional cast of the assessment of overall human "perfection").

¹¹⁹ See GRIFFIN, *supra* note __, at 58 ("Nor is there a *single* right balance [of objective goods]. The right balance is very likely to vary from person to person."); HURKA, *supra* note __, at 97-98 (same, with respect to account of human perfection rather than welfare).

implementation. But the accuracy of that procedure (more precisely, its accuracy, if perfectly performed) is swamped by the disadvantages just elaborated. So the proper comparison, at the level of decision procedures, is between CBA (or some refinement thereof) and other procedures that are *also* inaccurate, in various ways, but like CBA (or some refinement) can economize on decision costs, error rates, and transparency.

Specifically, CBA and other multidimensional procedures (QUALY-based assessment and direct multidimensional assessment) will generally be more accurate in tracking overall well-being than nonaggregative procedures, and than unidimensional aggregative procedures. The term “nonaggregative” is a residual, catch-all category. By this we mean any procedure that does not seek to determine (or to approximate) the aggregate effect of the project with respect to one or more (objective or hedonic) constituents of well-being, or prerequisites for well-being, or proxies for these. Good examples of nonaggregative procedures actually used by, or proposed for agencies, include the following: using a de minimus risk threshold for certain types of health, safety, or environmental risk (that is, banning any product that poses more than a de minimus risk of the given type, for example, a cancer risk);¹²⁰ subjecting proposed regulatory requirements to the constraint that they be “technologically feasible”;¹²¹ permitting or proscribing harmful activity depending upon whether the activity is “customary.”¹²²

Unidimensional, aggregative procedures pick out *one* (objective or hedonic) constituent of well-being, or prerequisite of well-being, or proxy for these, and enjoin agencies to maximize along this *one* dimension. The classic example, here, is so-called “Arisk-risk” analysis¹²³:

¹²⁰ Cf. LESTER LAVE, THE STRATEGY OF SOCIAL REGULATION 11-13 (1981) (describing no-risk approach). Agencies now more commonly employ a de minimus risk threshold than a true no-risk approach.

¹²¹ More precisely, this constraint is not a full decision-procedure itself, but will lend a nonaggregative element to whatever procedure is employed. On the feasibility constraint, see *id.* at 14-15.

¹²² Most familiarly used by courts in medical malpractice cases.

¹²³ See, e.g., LAVE, *supra* note __, at 15-17; Cass Sunstein, *Health-Health Tradeoffs*, 63 U. CHI. L. REV. 1533 (1996); W. Kip Viscusi, *Economic Foundations of the Current Regulatory Reform Efforts*, 10 J. ECON. PERSPECTIVES 119, 129-31 (1996).

in effect, “risk-risk” analysis tells the agency to compare the total number of premature deaths in the project world, with the total number of premature deaths in the status quo, and pick the world with the smaller number. The relevant dimension of well-being tracked by risk-risk analysis is longevity. One could imagine analogous procedures for any constituent of well-being, or any prerequisite: for example, maximizing the preservation of the wilderness, or the cleanliness of water, or the health of the population, or its educational opportunities, or the quality-adjusted size of the housing stock.

Finally, multidimensional procedures seek to track the aggregate effect of the project with respect to *more than one* dimension of well-being. These are best described by example, since in practice there are only three: CBA and refinements, direct multidimensional assessment, and QUALY-based assessment. The idea behind CBA and refinements is to reduce the project’s overall effect, on each person, to a single dollar amount (the CV, defined as welfare equivalent or as WTP/WTA) and then to aggregate. By contrast, direct multidimensional assessment instructs the analyst to calculate the aggregate effect of the project along each of several dimensions, and then to use either predefined quantitative tradeoff rates, or qualitative judgments, to compare aggregate project gains along the dimensions where its overall effect is positive, with aggregate project losses along the dimensions where its overall effect is negative.¹²⁴ Finally, QUALY-based assessment, a tool which is widely used by health economists and, to a lesser extent, by agencies, to evaluate health and risk-related projects, looks not merely to sheer longevity but to the *quality* of the life-years saved (or lost) by regulatory intervention. Information garnered from questionnaires is used to discount life years, relative to a baseline of perfect health; for example, the project of funding a medical program that will enable 100 beneficiaries to live, on average, 10 more years of life, but in a state of considerable pain, might be assessed as producing not 1000 life-years but $1000 \cdot .8 = 800$ *quality-adjusted* life years, with .8 as the discount factor for that kind of pain.¹²⁵

¹²⁴ See *infra* text accompanying notes ___ (describing different kinds of multidimensional assessment).

¹²⁵ See Robert Fabian, *The Quality Approach* in VALUING HEALTH FOR POLICY: AN ECONOMIC APPROACH 118 (George Tolley et al. eds., 1994).

Why should CBA, QUALY-based assessment, and direct multidimensional assessment be more accurate than unidimensional procedures or nonaggregative procedures? Overall welfare, itself, is multidimensional: this is directly true of an interpersonal construct like Griffin's (which incorporates multiple criteria of objective value or hedonic tone) and, indirectly, true of a desire-based construct like Harsanyi's (since persons typically develop preferences over world states, whether ordinary preferences or "extended preferences," in virtue, in part, of the multiple objective values or hedonic tones that they realize or believe they realize in those states).¹²⁶ Further, agency projects typically have multidimensional impacts, not just an impact upon the one dimension measured by the unidimensional procedure. Consider, as a paradigm, the dimension of longevity measured by straight "risk-risk" analysis. Agency projects that increase or decrease the total number of premature deaths, relative to the status quo, will typically accomplish more than just that with respect to welfare: they will also increase the stock or quality of goods, or the convenience of travel, or air quality, or electricity prices, or whatever. Such effects are completely missed by risk-risk analysis.¹²⁷ By contrast, a multidimensional procedure such as CBA will capture both the effect of the project on longevity, and its effect on other welfare-relevant dimensions. For example, a regulation that makes a product marginally safer, but significantly less enjoyable, or useful, or convenient will have a negative sum-of-CVs under CBA. This leads us to guess that risk-risk analysis will generally be less accurate, across the totality of agency choice situations, than CBA and other multidimensional procedures. At a minimum, and less ambitiously, we can say this: (1) in choice situations where the project has a relatively small effect on the dimension of longevity, such that further this effect is swamped by countervailing effects on other dimensions, risk-risk analysis will misclassify the project; (2) such choice situations constitute a significant fraction of the totality of agency choice situations, since longevity is neither the sole component of welfare, nor one that takes lexical priority over

¹²⁶ See, e.g., Griffin, *supra* note ____.

¹²⁷ Project gains along even a relatively unimportant dimension of human well-being, if large enough, can outweigh the welfare cost of premature death. See Alastair Norcross, *Comparing Harms: Headaches and Human Lives*, 26 PHIL. & PUB. AFF. 135 (1997).

others. Risk-risk analysis is inconsistent with art museums, public gardens, national holidays, and many other public projects. In this sense, risk-risk analysis embodies a significant procedural disadvantage, with respect to its accuracy in tracking overall welfare; and the same is true, *mutatis mutandis*, for other unidimensional procedures.¹²⁸

As for nonaggregative procedures: these typically capture some (purportedly) relevant aspect of agency choices other than the effect on overall well-being. For example, the *de minimus* risk threshold employed by some agencies for certain carcinogen risks seeks to prevent producers from wrongly or unfairly imposing cancer risks on unwitting consumers or workers; as a matter of overall well-being, however, a few more premature deaths from cancer might easily be justified by countervailing benefits, either the forestalling of premature deaths from other sources, or benefits unrelated to longevity. Similarly, the requirement that otherwise-justified agency choices be “technologically feasible” at best reflects a concern for the job security of workers in regulated industries, such that a technologically infeasible requirement (which would close down the industry) will not be imposed even if the benefits to consumers and citizens from doing so would outweigh the benefits from the industry’s continued existence. As far as we can see, the only standard nonaggregative procedure plausibly defended on welfarist grounds is the procedure of looking to social norms B given the possible efficiency of certain norms. Even here, however, a welfarist defense would need to be highly qualified and limited in scope. Robert Ellickson, the leading defender of the view that norms are efficient, claims that “members of a close-knit group develop and maintain norms whose content serves to maximize the aggregate welfare that members obtain in their workaday affairs with one another.”¹²⁹ Even if Ellickson’s claim is true (and one of us has argued elsewhere that it

¹²⁸ A more sophisticated approach would be to assign different unidimensional procedures to different agencies or statutory schemes. But multiple dimensions, other than the dimension tracked by the assigned procedure, might still be implicated by each agency’s choices. *Cf.* Richard J. Pierce, Jr., *The Role of Constitutional and Political Theory in Administrative Law*, 64 TEX. L. REV. 469, 473-81 (1985) (describing how statutes standardly authorize or require agencies to consider multiple values).

¹²⁹ ROBERT ELICKSON, *ORDER WITHOUT LAW: HOW NEIGHBORS SETTLE DISPUTES* 167 (1991) (emphasis omitted).

may not be),¹³⁰ a norm-based decision procedure for agencies will have a haphazard connection to overall welfare where the norm-generating group is *not* close knit, where “workaday” matters are not at stake, or where group outsiders are significantly affected by the agency’s choice.

To be sure, unidimensional and nonaggregative procedures will typically be cheaper, more transparent, and easier to implement correctly than CBA and other multidimensional procedures.¹³¹ But error costs and shirking costs can plausibly be held down with institutional mechanisms, for example the oversight of specialized cost-benefit bureaus such as OMB or the Congressional Budget Office.¹³² Assume, at the limit, that the additional decisional costs of the multidimensional procedure, relative to a unidimensional or nonaggregative procedure, is purely a matter of direct costs (the costs of information-gathering and processing by nonshirking, epistemically reliable agents). Then, for choice situations where the welfare effect of the project becomes sufficiently large, the multidimensional procedure will always be welfare-justified (assuming it’s more accurate than the alternative procedures), notwithstanding its additional decisional costs.

Turn now to a comparison of CBA and other multidimensional procedures. Consider, first, the procedure we have termed “direct multidimensional assessment.” This term actually names a family of related procedures. One variant of direct multidimensional assessment is fairly quantitative. On this variant, agencies are instructed to calculate or approximate the aggregate effect of the project along each of several predefined dimensions D_1, D_2, \dots, D_n , and then to use predefined tradeoff rates (one for each D_i - D_j combination, with $i < j$) to compare aggregate project gains along one dimension with aggregate project losses along another. This is the kind of procedure suggested by the remarks of Thomas Scanlon.

¹³⁰ Eric A. Posner, *Law, Economics, and Inefficient Norms*, 144 U. PA. L. REV. 1697 (1996).

¹³¹ Or at least arguably so; it is not clear whether procedures that incorporate the elastic criteria of “customary practice” and “technological feasibility” really have these virtues.

¹³² On the President’s ability to monitor agencies, including through the use of oversight bureaus, see, e.g., Terry M. Moe & Scott Wilson, *Presidents and the Politics of Structure*, LAW & CONTEMP. PROBS., Spring 1994.

Supposing, however, that we have formulated a conception of individual well-being appropriate for the purposes of moral and political argument, delimiting those personal interests that give rise to important claims on us and our shared institutions, there remains the further question of how institutional responses to these interests are to be measured and how individual distributive shares are to be compared . . . I will refer to such an answer . . . as an *index*. An index need not, like Rawls', consist simply of exchangeable goods and institutional prerogatives. It might refer as well to levels of development of personal capacities, as Sen has suggested, or even to states of consciousness. The avoidance of chronic physical pain, for example, might be one component in an index of well-being.
...

[A]n index of well-being is something that will be used by individuals, including legislators and other officials, in assessing institutional contributions to individual welfare.¹³³

A more qualitative version of direct multidimensional assessment would define the dimensions along which aggregate project impacts were to be assessed, but not the tradeoff rates. A yet more qualitative version would define neither: agencies would simply be instructed to assess (and report upon) aggregate impacts along "relevant" dimensions, and to use their judgment in making tradeoffs. Cass Sunstein and Richard Pildes suggest that agencies might engage in a variant of direct multidimensional assessment closer to the qualitative end of the spectrum.

We do not do well if we see such diverse goods as greater employment, protection of endangered species, lower prices, distributional effects, and cleaner air along a single [cost-benefit] metric, one that erases the qualitative differences among these goods. At least in principle, it would be better to have a disaggregated system for assessing the qualitatively different effects of regulatory impositions. . .

¹³³ Scanlon, *supra* note __, at 41.

Through considerations of this sort, we might be able to make some progress toward reform of existing cost-benefit analyses. Through regulatory-impact analyses, people should be allowed to see the diverse effects of regulation for themselves, and to make judgments based on an understanding of the qualitative differences.¹³⁴

Indeed, as we described in the Introduction, direct multidimensional assessment of the more qualitative kind is a procedure that agencies regularly employ in lieu of CBA.

Direct multidimensional assessment, of whatever variant, avoids the endowment-dependence characteristic of CBA. To see the point most simply: imagine a project that decreases the smog density over one city by 1 density unit, and increases the concentration of putrid airborne particles over another city with an identical population by 1 smell unit, such that, further, the interpersonal welfare tradeoff between a smoggy and a smelly atmosphere is 1 to 1. Then direct multidimensional assessment will accurately characterize the project as neither better nor worse than the status quo. If the population of each city is 1 million, direct multidimensional assessment will count the project as a decrease of 1 million along the smog dimension, and an increase of 1 million along the smell dimension, which counterbalance each other, given the 1-to-1 tradeoff between the dimensions. By contrast, if the population of the first city is richer than the population of the second, CBA will, inaccurately, characterize the project as a welfare improvement.¹³⁵

On the other hand, direct multidimensional assessment has its own disadvantages. First, assume that the correct interpersonal construct (as per Griffin) does indeed look solely to the effect of projects upon objective dimensions of value or hedonic tone. Even so,

¹³⁴ Richard H. Pildes & Cass R. Sunstein, *Reinventing the Regulatory State*, 62 U. CHI. L. REV. 1, 65 (1995).

¹³⁵ See Copp, *supra* note __, at 75 (“If Rachel and Paula are equally afflicted with asthma, then the clean air project might be of equal medical benefit to them, and this fact could be expressed in terms of the equal breathing efficiency they can expect if the project is implemented. Yet if Rachel is rich while Paula is poor, she may be willing to pay more than Paula for this benefit . . .”).

the number of such dimensions (if the construct is reasonable) will need to be large, which in turn tends to make the procedure of direct multidimensional assessment either less accurate, or less reliable, transparent, and cheap. To see *this* point, imagine the regulatory project of banning the sale of a certain kind of recreational boat, which poses a low but non-minimal and unavoidable death risk to users. The quantitative variant of multidimensional assessment might have a predefined dimension, "Recreational Value," such that 1-user-hour of that value is worth some fractional year less of longevity. But surely *that* objective assessment is much too crude; it implies, implausibly, that 1-user-hour with the boat has the very same welfare impact as 1-user-hour playing chess or 1-user-hour strolling in a public park. So the predefined dimensions might be made more fine grained. We might have a dimension for "Boating Recreation" B although even that would involve inaccuracies B or, finally, "Boating Recreation/Fishing" and "Boating Recreation/Sailing." But as the number of predefined dimensions (and tradeoff rates between them) increase, the decisional demands on the oversight bureau in maintaining this schema in accurate shape become large. This problem is avoided by the version of multidimensional assessment that does not predefine dimensions or tradeoffs; on the other hand, the absence of predefined dimensions or tradeoffs makes it much more difficult for the bureau, the legislature, or the public to monitor agencies.¹³⁶

The problem with direct multidimensional assessment becomes deeper if the right construct for interpersonal comparisons is *not* just what Griffin proposes B if it *does* incorporate some information about the strength of persons' desires, independent of the satisfaction of objective-value or hedonic criteria. Imagine that there are two kinds of

¹³⁶CBA avoids the problem sketched here because its basic building block is the individual CV, not units of overall value or hedonic tone. Decisionmakers can draw a demand curve for the recreational boat market without specifying, or fully specifying, what it is about the boat that makes it good. If CVs are defined in terms of WTP/WTA, then the relevant demand curve is simply the observed, market demand curve. If CVs are defined in terms of welfare equivalents, then decisionmakers will need to supplement market data B for example, they may need to provide boat buyers with fuller information, and determine how much money makes self-regarding buyers indifferent between having and not-having the fully described boat B but a complete specification of welfare dimensions may still be unnecessary.

boats, which are the same in objective and hedonic goodness, and yet consumers just like the second kind more. There is nothing about the second boat that justifies the consumers' stronger preference for it; the preference just is stronger, and remains so under full information. (Is this impossible? Think of preferences over flavors of ice cream.) So the demand curve for the second kind of boat will have a larger area under it than the demand curve for the first kind. CBA will properly reflect that difference, while no variant of direct multidimensional assessment can.

A final problem with direct multidimensional assessment concerns the individualization of tradeoff rates. Aggregate recreational value, for the boating market, is being compared with the aggregate increase in longevity that would result from banning the boat. But, as we remarked earlier, even on the Griffin-type construct for making interpersonal comparisons, the true tradeoff rate will vary from person to person. For A, 1 hour of Recreational Value/Fishing is worth 2 hours less of longevity. For B, 1 hour of Recreational Value/Fishing is worth 3 hours less. The demand curve or estimated demand curve, in dollars, for the boat will reflect the mix of recreation/longevity tradeoff rates among the population of consumers. By contrast (at least where predefined nationwide or regional tradeoffs rates are specified, which would presumably be the case with the quantitative variant of direct multidimensional assessment), this welfare-relevant information about the boating market would be lost.

What about QUALY-based assessment? Sunstein and Pildes identify this procedure as another possible alternative to conventional CBA.¹³⁷ But note that QUALY-based assessment is standardly used, and conceptualized, as a *cost-effectiveness* tool for evaluating health and risk regulation and expenditure.¹³⁸ That is: given a fixed dollar budget (which could be a budget for direct governmental expenditures, or a "regulatory budget" for compliance costs) the agency chooses the policy that maximizes QUALYs. Another way to put the point is that QUALY-based assessment has limited scope: the project, the status quo, and all other options being compared must be identical except on the dimension of health and risk. Imagine that the status quo involves

¹³⁷ See Pildes & Sunstein, *supra* note __, at 83-85.

¹³⁸ See Fabian, *supra* note __. See generally LAVE, *supra* note __, at 19-23 (discussing cost-effectiveness analysis).

a baseline governmental expenditure of \$100 million for a saving of 200 quality-adjusted-life-years, and the project involves a governmental expenditure of a different amount, say \$500 million, for a saving of 1100 quality-adjusted-life-years. Standard QUALY analysis has nothing to say about this comparison.

In theory, the standard QUALY procedure could be reconceptualized as a general multidimensional tool: one would simply calculate the total number of quality-adjusted-life-years in both the project world and the status quo year, with quality adjustments not just for health effects but for *any* kind of effect on welfare.¹³⁹ The problem, here, is translating non-health effects into quality discounts or premiums. By what factor does the eating of an ice cream cone, once, change the quality of the year in which it is consumed? By what factor does the experience of walking in an old-growth forest? A faster daily commute home? Better visibility over your home? To be sure, CBA also involves quantifying these welfare effects, but we see a number of obvious reasons why the translation of ice cream cones, wilderness walks, faster commutes, and so on into dollar CVs would be performed more accurately, cheaply, and transparently, by agencies, than their translation into discounts or premiums for life-years. First, many goods, or their welfare analogues, are traded on markets; and the market price of a good is at least some evidence, perhaps strong evidence, of its CV for a given person, depending on how the concept of CV is defined. Second, the very fact that a person routinely trades on markets will make it easier for her to conceptualize welfare impacts in dollar terms; however difficult the respondent to a CV-survey¹⁴⁰ might find it to value ice cream cones, wilderness walks, etc., in dollar terms, we would expect that the respondent to a QUALY-survey would find

¹³⁹ For example, if the \$400 million difference in the above hypothetical would be spent on a 1-year arts program in the status quo, such that each of 100,000 viewers would experience a .005 increase in the quality of her year, the \$400 million difference could be translated into an increment of 500 quality-adjusted-life years for the status quo, relative to the project. Alternately, as some suggest, *see* Fabian, *supra* note __, a standard-quality year could be monetized, along with other non-health benefits; but this is just then a kind of CBA.

¹⁴⁰ On the use of so-called "contingent-valuation" surveys to determine CVs, see MITCHELL & CARSON, *supra* note __; CONTINGENT VALUATION: A CRITICAL ASSESSMENT (Jerry A. Hausman ed., 1993); Symposium, *Contingent Valuation*, J. ECON. PERSP., Fall 1994, at 3.

the valuation yet more difficult, and the answers to QUALY-surveys (for these non-health effects) yet less reliable.¹⁴¹

* * *

The comparison between CBA and other multidimensional procedures is a complex one. We have no doubt that, for some choice situations (for example, for a choice between two large health-related policies that have identical non-health impacts), either QUALY-based assessment or direct multidimensional assessment will turn out to be the welfare-maximizing procedure. But the case for a wholesale replacement of CBA has not been made and, we think, cannot be. Our analysis suggests that QUALY based assessment is best limited to use as a cost-effectiveness tool, and that agencies are welfare-justified in engaging in direct multidimensional assessment in lieu of CBA only where wealth effects are likely to be large (for example, where the winners or losers are significantly richer or poorer than the losers or winners). CBA has flaws, but so do its alternatives B including flaws with respect to accuracy in tracking overall well-being, something that has been a standard complaint about CBA.¹⁴²

F. Non-welfarist considerations

Sometimes it is objected that CBA fails to reflect nonwelfarist considerations, for example, considerations of fair distribution;¹⁴³ and sometimes it is said, in response to this objection, that to talk of fairness

¹⁴¹ Finally, and on a different note, it should be observed that QUALY-based assessment is not a perfect welfarist procedure, quite apart from these issues of transparency, reliability, and decisional cost. The addition of one standard-quality year to the life of a long-lived person arguably has less impact on overall well-being than the addition of a standard-quality year to the life of someone with fewer total years to her life. Longevity (quality adjusted or not) has declining marginal utility, just like wealth; and thus QUALY-based assessment has its own analogue to the problem of endowment-dependence that afflicts CBA.

¹⁴² Our defense of CBA is not meant to be a criticism of market-based alternatives, where those are appropriate. The great advantage of the latter is that they do not put great strain on the ability of agencies to gather and evaluate information. CBA is not a rejection of command-and-control regulation and not a solution to all its problems. See Cass R. Sunstein, *Free Markets and Social Justice* 375-76 (1997).

¹⁴³ See, e.g., Lester B. Lave, *Benefit-Cost Analysis: Do the Benefits Exceed the Costs?*, in *RISKS, COSTS, AND LIVES SAVED* 104, 114-15 (Robert W. Hahn ed., 1996).

is simply to express a preference, and that the preference for fairness (like any other) would indeed figure in a proper accounting of costs and benefits.¹⁴⁴ In our view, both the objection and the response are misconceived.

Start with the response. The response is misconceived on two counts. First, it embodies a *metaethical* confusion, about the nature of moral facts. It is a moral fact that, everything else equal, government projects that promote overall well-being are better than government projects that do not. It is also a moral fact that government should not deliberately murder an identified, innocent person, even if doing so increases overall well-being. This latter, moral fact does not reduce to preferences; even if virtually everyone preferred to murder the one innocent, doing so would still be wrong. Nor does the former moral fact: even if virtually everyone preferred, say, projects that benefitted an elite which had succeeded in brainwashing the population to believe the elites to be morally superior, it would still remain a moral fact that, as between a project P1 that improves overall well-being and a project P2 that benefits the elite, P1 is morally better. The metaethical notion that a moral fact reduces to a fact about what everyone, or the majority prefers, is an implausible version of conventionalism about ethics.¹⁴⁵

Second, the response conflates preference and well-being. Satisfying certain preferences *is* important, but this is not a foundational matter; it is not because, at bottom, moral claims are claims about preferences. Rather, it is true *within* ethics because, among other things, overall well-being is morally important, and certain preferences are constitutive of well-being. Which ones? That depends, as we have already explained, on your theory of well-being. The right theory is a desire-based theory; and, if “preference” is taken broadly to mean desire, i.e., a pro-attitude, then people are indeed better

¹⁴⁴ See generally Louis Kaplow and Steven Shavell, *Fairness Versus Welfare Economics in Normative Analysis of the Law* (unpublished manuscript 1998). Although Kaplow and Shavell’s views fall into this general category, their argument is complex and the reader should make no inferences about it from our discussion in the text.

¹⁴⁵ On ethical conventionalism, see, e.g., BRINK, *supra* note __, at 14-36. At best, moral facts reduce to facts about ideally informed preferences, which are still different from facts about what the majority or all of us actually prefer. For an argument that moral facts do reduce to facts about ideally informed preferences, see MICHAEL SMITH, *THE MORAL PROBLEM* (1994).

off to the extent that certain of their preferences are satisfied. The word “certain” deserves emphasis, here, because desire-based theories might require restricted desires, or informed desires, or even good desires. If a project will alleviate the misery of the rural poor, then the misanthrope who has the sadistic preference to see the rural poor suffer is arguably *not* better off if in fact they do suffer. Sadistic preferences might not be—indeed, plausibly are *not*—constitutive of well-being.¹⁴⁶ Similarly, if a project will alleviate the misery of the rural poor, and a concerned taxpayer (a) concludes that distributive justice requires helping them, and therefore (b) prefers a rural-poverty project, then this fairness-preference might not be—indeed, plausibly is *not*—constitutive of the taxpayer’s well-being.¹⁴⁷ Do we want to say that the rural project is supported both by considerations of fair distribution and, in addition, by the taxpayer’s well-being? That seems implausible. We do not believe that the fairness of a project increases its CV, insofar as persons prefer the project just because they judge it to be fair. An economist who disagrees will need to argue, more specifically than economists have done, about the nature of well-being; he will need to show why preferences for fairness are constitutive of well-being, just like preferences for concerts, widgets, and skiing. And even if he shows this, he will need to admit that the moral force of fairness is not exhausted by preference: a project might be, all things considered, morally wrong (say, distributively unjust) even though it does improve overall well-being.

So our welfarist defense of CBA is nested within a view of morality that is cognitivist¹⁴⁸ (not skeptical) at the metaethical level, and non-utilitarian at the substantive level. Again: overall well-being is morally relevant, not morally decisive. To claim moral *decisiveness* is to affirm utilitarianism, which famously leads to a variety of counterintuitive moral positions, for example, that killing one to save two is morally required, or that persons are obliged to abandon their personal pursuits if doing so would increase overall well-being.¹⁴⁹ The

¹⁴⁶ See, e.g., Harsanyi, *supra* note __, at 56.

¹⁴⁷ See *supra* text accompanying note __ (arguing for a restricted-desire account of welfare).

¹⁴⁸ On cognitivism, see Darwall, Gibbard & Railton, *supra* note __.

¹⁴⁹ See Bernard Williams, *A Critique of Utilitarianism*, in UTILITARIANISM: FOR AND AGAINST 77 (1963).

non-welfarist considerations that bear upon the worth of governmental projects may include some or all of the following: (1) deontological considerations, specifically, constraints against performing projects that maximize good consequences;¹⁵⁰ (2) egalitarian or “distributive” considerations, such as the effect of the project on the welfare of the least-well-off group, or on persons below the poverty line, or its effect on the equality of welfare or resources;¹⁵¹ (3) desert-based considerations, namely, the extent to which the project rewards the deserving and harms the culpable;¹⁵² (4) perfectionist considerations, such as the purported intrinsic good of preserving endangered species or ecosystems or, more generally, good consequences produced by the project quite apart from any effect on welfare.¹⁵³ We do not mean to commit ourselves to a particular view about the nonwelfarist component of morality; we simply mention these four as standardly-discussed possibilities.¹⁵⁴

CBA does not plausibly capture, and is not meant to capture, non-welfarist considerations. The objection that CBA fails to capture them is really no objection at all B any more than, say, the failure of a statistical measure of equality to track overall well-being (rather than equality) is an objection to the proposition that agencies should use that measure in assessing the distributional consequences of its projects. We do not conceptualize CBA as the exclusive choice procedure for government, but as one part of the overall set of procedures and institutions by which projects are ultimately approved, rejected, or amended. How nonwelfarist considerations *should* be captured B whether agencies should generally use CBA and non-welfarist procedures seriatim; or instead there should be a separation of welfarist and non-welfarist considerations between agencies; or instead non-

¹⁵⁰ See *supra* text accompanying note ___.

¹⁵¹ See *supra* text accompanying note ___.

¹⁵² See GEORGE SHER, DESERT (1987).

¹⁵³ See HURKA, *supra* note ___; HOLMES ROLSTON, III, ENVIRONMENTAL ETHICS: DUTIES TO AND VALUES IN THE NATURAL WORLD 126-59 (1988).

¹⁵⁴ A further point is that if people’s moral preferences are sufficiently widespread and uniform, then a project of maximizing social welfare reduces to a project of respecting people’s moral views. Cost-benefit analysis, or a procedure of maximizing social welfare, would be methodologically empty because parasitic on the philosophical project of ascertaining people’s morality.

welfarist considerations should uniquely be the concern of the legislature, the court system, or both. B is too complex an issue to be considered here.

CONCLUSION

CBA is a useful decision procedure and it should be routinely used by agencies. CBA appears to be superior to rival methodologies in enabling agencies to evaluate projects according to the extent that they contribute to overall well-being. It allows agencies to take into account all relevant influences on overall well-being, unlike simpler decision procedures like risk-risk; and it enables agencies to weigh the advantages and disadvantages of projects in a clear and systematic way, unlike more complex decision procedures. Because maximizing overall well-being is an important role of the government in all major political theories, and is consistent with widespread intuitions, it is a worthy goal of agency action. Finally, CBA plays the important political role of increasing regulatory transparency. The political branches can monitor agencies more easily when the agencies monetize the advantages and disadvantages of projects than when agencies use qualitative decision procedures.

However, CBA serves these useful purposes only under certain conditions, and agencies should take account of these conditions when evaluating projects. First, CBA must give way to important non-welfarist concerns, such as deontological rights. An agency should not approve a project that has a positive sum of CVs if it involves the unjustified sacrifice of an innocent. Second, CBA must give way when the endowments of affected people vary a great deal. Either the agency should attempt to adjust CVs in light of wealth differences, or it should abandon CBA in these circumstances in favor of a more complex decision procedure, or perhaps it should encourage the political branches to construct a deal that compensates the losers. Third, CBA must be adjusted to account for uninformed or distorted preferences. Agencies may be able to overcome these problems by informing individuals before determining individuals' CVs, or even by engaging in an imaginative reconstruction of their preferences. But in all these cases, the extent to which the agency should defer to people's stated preferences or not must depend on its competence, the importance of transparency, the effectiveness of political monitors, and so on.

Interestingly, agencies already seem to depart from textbook CBA in order to respond to these concerns. As the pesticide case study shows, agencies are sensitive to wealth differentials. EPA did not assume lower CVs for risk of death for migrant farm workers than for the average person. As the Grand Canyon case study shows, agencies are sensitive to information problems. However crude the technique may have seemed, the agency clearly understood that informed preferences were more relevant than uninformed preferences. As the lead-in-drinking-water case study shows, difficulties in monetizing benefits can be overcome. Of course, the departures from standard CBA methods create some uncertainty, and this uncertainty makes it more difficult to evaluate agency decisions. Nevertheless, compromise of some sort is sensible and unavoidable, and EPA's behavior shows that while CBA methods cannot be applied mechanically, agencies can use it to guide judgment, in a way that rationalizes and clarifies agency action.

The paper is a draft of an article that has been accepted for publication in the *Yale Law Journal*, Vol. 109, 1999. All rights reserved. Readers with comments should address them to:

Eric A. Posner
Professor of Law
University of Chicago Law School
1111 East 60th Street
Chicago, IL 60637

eric_posner@law.uchicago.edu

Chicago Working Papers in Law and Economics
(Second Series)

1. William M. Landes, Copyright Protection of Letters, Diaries and Other Unpublished Works: An Economic Approach (July 1991).
2. Richard A. Epstein, The Path to *The T.J. Hooper*: The Theory and History of Custom in the Law of Tort (August 1991).
3. Cass R. Sunstein, On Property and Constitutionalism (September 1991).
4. Richard A. Posner, Blackmail, Privacy, and Freedom of Contract (February 1992).
5. Randal C. Picker, Security Interests, Misbehavior, and Common Pools (February 1992).
6. Tomas J. Philipson & Richard A. Posner, Optimal Regulation of AIDS (April 1992).
7. Douglas G. Baird, Revisiting Auctions in Chapter 11 (April 1992).
8. William M. Landes, Sequential versus Unitary Trials: An Economic Analysis (July 1992).
9. William M. Landes & Richard A. Posner, The Influence of Economics on Law: A Quantitative Study (August 1992).
10. Alan O. Sykes, The Welfare Economics of Immigration Law: A Theoretical Survey With An Analysis of U.S. Policy (September 1992).
11. Douglas G. Baird, 1992 Katz Lecture: Reconstructing Contracts (November 1992).
12. Gary S. Becker, The Economic Way of Looking at Life (January 1993).
13. J. Mark Ramseyer, Credibly Committing to Efficiency Wages: Cotton Spinning Cartels in Imperial Japan (March 1993).
14. Cass R. Sunstein, Endogenous Preferences, Environmental Law (April 1993).
15. Richard A. Posner, What Do Judges and Justices Maximize? (The Same Thing Everyone Else Does) (April 1993).
16. Lucian Arye Bebchuk and Randal C. Picker, Bankruptcy Rules, Managerial Entrenchment, and Firm-Specific Human Capital (August 1993).
17. J. Mark Ramseyer, Explicit Reasons for Implicit Contracts: The Legal Logic to the Japanese Main Bank System (August 1993).
18. William M. Landes and Richard A. Posner, The Economics of

- Anticipatory Adjudication (September 1993).
19. Kenneth W. Dam, The Economic Underpinnings of Patent Law (September 1993).
 20. Alan O. Sykes, An Introduction to Regression Analysis (October 1993).
 21. Richard A. Epstein, The Ubiquity of the Benefit Principle (March 1994).
 22. Randal C. Picker, An Introduction to Game Theory and the Law (June 1994).
 23. William M. Landes, Counterclaims: An Economic Analysis (June 1994).
 24. J. Mark Ramseyer, The Market for Children: Evidence from Early Modern Japan (August 1994).
 25. Robert H. Gertner and Geoffrey P. Miller, Settlement Escrows (August 1994).
 26. Kenneth W. Dam, Some Economic Considerations in the Intellectual Property Protection of Software (August 1994).
 27. Cass R. Sunstein, Rules and Rulelessness, (October 1994).
 28. David Friedman, More Justice for Less Money: A Step Beyond *Cimino* (December 1994).
 29. Daniel Shaviro, Budget Deficits and the Intergenerational Distribution of Lifetime Consumption (January 1995).
 30. Douglas G. Baird, The Law and Economics of Contract Damages (February 1995).
 31. Daniel Kessler, Thomas Meites, and Geoffrey P. Miller, Explaining Deviations from the Fifty Percent Rule: A Multimodal Approach to the Selection of Cases for Litigation (March 1995).
 32. Geoffrey P. Miller, Das Kapital: Solvency Regulation of the American Business Enterprise (April 1995).
 33. Richard Craswell, Freedom of Contract (August 1995).
 34. J. Mark Ramseyer, Public Choice (November 1995).
 35. Kenneth W. Dam, Intellectual Property in an Age of Software and Biotechnology (November 1995).
 36. Cass R. Sunstein, Social Norms and Social Roles (January 1996).
 37. J. Mark Ramseyer and Eric B. Rasmusen, Judicial Independence in Civil Law Regimes: Econometrics from Japan (January 1996).
 38. Richard A. Epstein, Transaction Costs and Property Rights: Or Do Good Fences Make Good Neighbors? (March 1996).

39. Cass R. Sunstein, *The Cost-Benefit State* (May 1996).
40. William M. Landes and Richard A. Posner, *The Economics of Legal Disputes Over the Ownership of Works of Art and Other Collectibles* (July 1996).
41. John R. Lott, Jr. and David B. Mustard, *Crime, Deterrence, and Right-to-Carry Concealed Handguns* (August 1996).
42. Cass R. Sunstein, *Health-Health Tradeoffs* (September 1996).
43. Douglas G. Baird, *The Hidden Virtues of Chapter 11: An Overview of the Law and Economics of Financially Distressed Firms* (March 1997).
44. Richard A. Posner, *Community, Wealth, and Equality* (March 1997).
45. William M. Landes, *The Art of Law and Economics: An Autobiographical Essay* (March 1997).
46. Cass R. Sunstein, *Behavioral Analysis of Law* (April 1997).
47. John R. Lott, Jr. and Kermit Daniel, *Term Limits and Electoral Competitiveness: Evidence from California's State Legislative Races* (May 1997).
48. Randal C. Picker, *Simple Games in a Complex World: A Generative Approach to the Adoption of Norms* (June 1997).
49. Richard A. Epstein, *Contracts Small and Contracts Large: Contract Law through the Lens of Laissez-Faire* (August 1997).
50. Cass R. Sunstein, Daniel Kahneman, and David Schkade, *Assessing Punitive Damages (with Notes on Cognition and Valuation in Law)* (December 1997).
51. William M. Landes, Lawrence Lessig, and Michael E. Solimine, *Judicial Influence: A Citation Analysis of Federal Courts of Appeals Judges* (January 1998).
52. John R. Lott, Jr., *A Simple Explanation for Why Campaign Expenditures are Increasing: The Government is Getting Bigger* (February 1998).
53. Richard A. Posner, *Values and Consequences: An Introduction to Economic Analysis of Law* (March 1998).
54. Denise DiPasquale and Edward L. Glaeser, *Incentives and Social Capital: Are Homeowners Better Citizens?* (April 1998).
55. Christine Jolls, Cass R. Sunstein, and Richard Thaler, *A Behavioral Approach to Law and Economics* (May 1998).
56. John R. Lott, Jr., *Does a Helping Hand Put Others At Risk?: Affirmative Action, Police Departments, and Crime* (May 1998).

<http://www.law.uchicago.edu/Publications/Working/index.html>

57. Cass R. Sunstein and Edna Ullmann-Margalit, *Second-Order Decisions* (June 1998).
58. Jonathan M. Karpoff and John R. Lott, Jr., *Punitive Damages: Their Determinants, Effects on Firm Value, and the Impact of Supreme Court and Congressional Attempts to Limit Awards* (July 1998).
59. Kenneth W. Dam, *Self-Help in the Digital Jungle* (August 1998).
60. John R. Lott, Jr., *How Dramatically Did Women's Suffrage Change the Size and Scope of Government?* (September 1998)
61. Kevin A. Kordana and Eric A. Posner, *A Positive Theory of Chapter 11* (October 1998)
62. David A. Weisbach, *Line Drawing, Doctrine, and Efficiency in the Tax Law* (November 1998)
63. Jack L. Goldsmith and Eric A. Posner, *A Theory of Customary International Law* (November 1998)
64. John R. Lott, Jr., *Public Schooling, Indoctrination, and Totalitarianism* (December 1998)
65. Cass R. Sunstein, *Private Broadcasters and the Public Interest: Notes Toward A "Third Way"* (January 1999)
66. Richard A. Posner, *An Economic Approach to the Law of Evidence* (February 1999)
67. Yannis Bakos, Erik Brynjolfsson, Douglas Lichtman, *Shared Information Goods* (February 1999)
68. Kenneth W. Dam, *Intellectual Property and the Academic Enterprise* (February 1999)
69. Gertrud M. Fremling and Richard A. Posner, *Status Signaling and the Law, with Particular Application to Sexual Harassment* (March 1999)
70. Cass R. Sunstein, *Must Formalism Be Defended Empirically?* (March 1999)
71. Jonathan M. Karpoff, John R. Lott, Jr., and Graeme Rankine, *Environmental Violations, Legal Penalties, and Reputation Costs* (March 1999)
72. Matthew D. Adler and Eric A. Posner, *Rethinking Cost-Benefit Analysis* (April 1999)