

Southern Methodist University

SMU Scholar

Faculty Journal Articles and Book Chapters

Faculty Scholarship

2010

The Endogeneity Problem in Cost-Benefit Analysis

Gregory S. Crespi

Southern Methodist University, Dedman School of Law

Recommended Citation

Gregory Scott Crespi, *The Endogeneity Problem in Cost-Benefit Analysis*, 8 *Geo. J.L. & Pub. Pol'y* 91 (2010)

This document is brought to you for free and open access by the Faculty Scholarship at SMU Scholar. It has been accepted for inclusion in Faculty Journal Articles and Book Chapters by an authorized administrator of SMU Scholar. For more information, please visit <http://digitalrepository.smu.edu>.

The Endogeneity Problem In Cost-Benefit Analysis

GREGORY SCOTT CRESPI*

INTRODUCTION	92
I. THE ENDOGENOUS PREFERENCES PROBLEM	97
A. <i>A General Introduction</i>	97
B. <i>Prior Efforts to Address the Problem of Endogenous Preferences in Cost-Benefit Analysis</i>	101
1. The Becker “Extended Utility Function” Approach	101
2. Other Commentary Regarding Incorporating Endogenous Preferences in Cost-Benefit Analysis	104
(i) <i>Cass Sunstein’s Views</i>	104
(ii) <i>Kenneth Dau-Schmidt’s Views</i>	107
(iii) <i>Samuel Bowles’ Views</i>	108
C. <i>My Approach to Making Valuations in Cost-Benefit Analysis Under Endogenous Preferences Circumstances</i>	109
1. Practical Considerations in Applying an Endogenous Preferences Valuation Framework	110
(i) <i>Under Some Circumstances the Choice Among Alternative Preference Structures Is Mooted</i>	110
(ii) <i>Preference Structure Estimation Difficulties May Limit the Ability to Implement the Theoretically Appropriate Valuation Approach</i>	111
(a) General Preference Structure Estimation Difficulties	111
(b) Prospective and Retrospective Cost-Benefit Analyses Present Different Kinds of Problems for Determining the Applicable Preference Structures	112
2. A Proposed Methodology for Combining Valuations Obtained Through the Use of Different Preference Structures	113

* Professor of Law, Dedman School of Law, Southern Methodist University. J.D., Yale Law School, Ph.D., University of Iowa. © 2010, Gregory Scott Crespi.

3.	Recognizing Practical Limitations on an Endogenous Preferences Valuation Methodology	116
D.	<i>Conclusions Regarding Incorporating Endogeneity of Preferences</i>	117
II.	THE ENDOGENEITY OF IDENTITY PROBLEM	118
A.	<i>A General Introduction</i>	118
B.	<i>The Problem of Person-Altering Consequences</i>	121
C.	<i>The Implications of Person-Altering Consequences for Cost-Benefit Analysis</i>	127
1.	Applying Standard Cost-Benefit Valuation Techniques to Assess Person-Altering Consequences	127
2.	Consideration of Two Alternative Cost-Benefit Approaches for Valuing Person-Altering Consequences	133
3.	Developing Alternative Criteria for Assessing the Impacts of Person-Altering Consequences	139
D.	<i>Conclusions Regarding the Problem of Endogeneity of Identity</i>	142
III.	CONCLUSION	144

INTRODUCTION

Cost-benefit analysis is the most important policy evaluation technique now used in American public sector decision-making.¹ This approach is claimed to be effective for helping policy makers identify measures that will broadly advance social welfare,² and as a means of countering the efforts of special interest groups, which often resist the implementation of such measures.³ Cost-benefit analysis plays a particularly important role in determining federal environmental standards as well as in other federal administrative rulemaking.⁴

1. See generally CASS R. SUNSTEIN, *THE COST-BENEFIT STATE: THE FUTURE OF REGULATORY PROTECTION* (Am. Bar Ass'n 2002) [hereinafter SUNSTEIN 2002]. Some recent writers use the inverted phrase "benefit-cost analysis" to describe this form of analysis, but in this Article I will consistently use the traditional label.

2. "[I]n cost-benefit analysis we are concerned with the economy as a whole, with the welfare of a defined society, and not any smaller part of it." E.J. MISHAN, *COST-BENEFIT ANALYSIS*, at x (Praeger Publishers 1976).

3. See SUNSTEIN 2002, *supra* note 1, at 26–28.

4. This method of policy evaluation was used to some extent in regulatory reviews under the Nixon, Ford, and Carter Administrations. President Nixon first created a "Quality of Life Review" that gave the Office of Management and Budget limited regulatory review authority. President Ford then required several federal agencies to also provide inflation impact statements for review by the Council on Wage

The prominent legal scholar Cass Sunstein, the current Director of the Office of Management and Budget's Office of Information and Regulatory Affairs, has gone so far as to claim that "[the] American government is becoming a cost-benefit state,"⁵ and to endorse this development as helping to rationalize government decision-making and insulate it from the pressure of special interest groups, as well as increasing regulatory transparency and public accountability.⁶

While Sunstein's descriptive claim regarding the extensive use of the cost-benefit methodology appears to be well founded,⁷ his endorsement of this trend is more problematic.⁸ In particular, there is considerable controversy regarding whether cost-benefit analysis provides an adequately inclusive and unbiased means of assessing policies and programs.⁹ The literature on this topic is

and Price Stability. President Carter took another significant step towards institutionalizing cost-benefit analysis as an aspect of federal rulemaking when he established the Regulatory Analysis Review Group and issued Executive Order 12,044, which required economic impact statements for all proposed rules having an overall impact of more than \$100 million. See generally Murray Weidenbaum, *Regulatory Process Reform from Ford to Clinton*, 20 REG. (1997), available at <http://www.cato.org/pubs/regulation/regv20n1/reg20n1a.html>; THOMAS O. MCGARTY, *REINVENTING RATIONALITY: THE ROLE OF REGULATORY ANALYSIS IN THE FEDERAL BUREAUCRACY* (Cambridge Univ. Press 1991).

The significance of cost-benefit analysis for federal rulemaking was greatly enhanced by Executive Order 12,291, issued in 1981 by President Reagan. See Exec. Order No. 12,291, 46 Fed. Reg. 13,193 (Feb. 17, 1981). Executive Order 12,291 required many proposed executive branch regulatory initiatives to be accompanied by a Regulatory Impact Analysis containing an extensive cost-benefit analysis of the proposal that had to be submitted to and approved by the Office of Information and Regulatory Affairs, a new office created within the Office of Management and Budget, before the proposed regulation could become effective. The Clinton Administration later replaced Executive Order 12,291 with a new Executive Order 12866 which made some minor changes in wording and procedures, but which retained the substance of the cost-benefit analysis requirement for major rulemaking initiatives. See Exec. Order No. 12,291, 58 Fed. Reg. 51,735 (Sept. 30, 1993). The Bush Administration also continued this policy.

Congress has also enacted numerous statutes in recent years requiring federal agencies to perform cost-benefit analyses in connection with their rulemaking efforts. See Matthew D. Adler & Eric A. Posner, *Rethinking Cost-Benefit Analysis*, 109 YALE L.J. 165, 167 (1999) [hereinafter Adler & Posner 1999]. Prospective cost-benefit analyses of rulemaking initiatives and subsequent Office of Information and Regulatory Affairs review (and, upon occasion, also judicial review) therefore now appear to be entrenched as a significant feature of the federal regulatory process. Similarly, cost-benefit analyses are now also utilized to a lesser extent by many state governmental agencies for similar purposes. See SUNSTEIN 2002, *supra* note 1, at 26–28. See generally Robert W. Hahn, *State and Federal Regulatory Reform: A Comparative Analysis*, 29 J. LEGAL STUD. 873 (2000).

5. SUNSTEIN 2002, *supra* note 1, at ix.

6. *Id.* at 26–28.

7. See Matthew D. Adler & Eric A. Posner, *Implementing Cost-Benefit Analysis When Preferences Are Distorted*, 29 J. LEGAL STUD. 1105, 1105 (2000); Robert W. Hahn & Paul C. Tetlock, *Has Economic Analysis Improved Regulatory Decisions?*, 22 J. ECON. PERSP. 67, 68 (2008). But see Amy Sinden, *Cass Sunstein's Cost-Benefit Lite: Economics for Liberals*, 29 COLUM. J. ENVTL. L. 191 (2004) (criticizing Sunstein's claim as to the pervasiveness of cost-benefit analysis).

8. See Douglas A. Kysar, *It Might Have Been: Risk, Precaution and Opportunity Costs*, 22 J. LAND USE & ENVTL. L. REV. 1, 5–6 (2006) ("The basic superiority of CBA [cost-benefit analysis] as a policy tool for risk regulation . . . is no longer seriously doubted . . . [e]xcept, of course . . . by serious, thoughtful observers of the administrative state."). See generally Sinden, *supra* note 7.

9. See Sinden, *supra* note 7, at 201 (citing voluminous literature offering criticisms of cost-benefit analysis); Adler & Posner 1999, *supra* note 4, at 167 ("The reputation of cost-benefit analysis . . . among American academics has never been as poor as it is today."). See generally Frank Ackerman &

extensive and includes a significant number of articles that sharply criticize this approach.¹⁰

The essential feature of cost-benefit analysis is its embrace of the willingness to pay valuation criterion. All impacts of the policies under consideration are measured, to the extent feasible, by the affected persons' willingness to pay to obtain the resulting benefits or to avoid the resulting burdens. Those benefit and cost measures are then aggregated into bottom-line economic efficiency assessments that do not address distributional considerations or rights-respecting limitations.¹¹ Such additional considerations and limitations are taken into account separately in the decision-making process, if at all.

The numerous criticisms of this methodology can be roughly, but usefully, classified as either "external" or "internal."¹² The external critiques largely reject the approach altogether, generally emphasizing the measurement problem presented by what those critics regard as the fundamental incommensurability of policy impacts of different character. The critics commonly conclude that cost-benefit analysis is fatally flawed because of the impossibility of meaningfully measuring diverse impacts by a unitary monetary metric. These impacts range from purely financial consequences to loss of life itself, and also include difficult to quantify effects such as the empathetic sentiments aroused in support of those persons more directly affected by a policy.¹³ These external critiques also commonly focus upon the well-known shortcomings of the use of economic efficiency as a normative standard.¹⁴

Lisa Heinzerling, *Pricing the Priceless: Cost-Benefit Analysis of Environmental Protection*, 150 U. PA. L. REV. 1553 (2002) (presenting a highly critical view of cost-benefit analysis); Kysar, *supra* note 8 (presenting a highly critical view of cost-benefit analysis).

10. See, e.g., Adler & Posner 1999, *supra* note 4; Ackerman & Heinzerling, *supra* note 9; Kysar, *supra* note 8.

11. Stated in more theoretical terms, cost-benefit analysis is an attempt to determine the Kaldor-Hicks efficiency consequences of a policy so that this information can guide the decision whether or not to implement that policy. A policy will constitute a Kaldor-Hicks improvement—a move towards Kaldor-Hicks efficiency—if the total benefits of the policy exceed its total costs, with both benefits and costs measured by the willingness to pay of the affected persons to obtain those benefits or avoid those costs. See RICHARD A. POSNER, *ECONOMIC ANALYSIS OF LAW* 13 (6th ed. 2003). The Kaldor-Hicks criterion is the usual measure of efficiency utilized by economists. *Id.* There are several seminal articles and books that developed the Kaldor-Hicks efficiency concepts. See, e.g., John R. Hicks, *The Foundations of Welfare Economics*, 49 *ECON. J.* 696 (1939); Nicholas Kaldor, *Welfare Propositions of Economics and Interpersonal Comparisons of Utility*, 49 *ECON. J.* 549 (1939); ANTHONY E. BOARDMAN ET AL., *COST-BENEFIT ANALYSIS AND PRACTICE* (Prentice Hall 2000) (1996); MISHAN, *supra* note 2.

12. Sinden, *supra* note 7, at 202.

13. *Id.* There is extensive literature advancing external critiques of the cost-benefit methodology based on the incommensurability problem. See, e.g., Ackerman & Heinzerling, *supra* note 9, at 1553, 1563–64; Henry S. Richardson, *The Stupidity of the Cost-Benefit Standard*, 29 *J. LEGAL STUD.* 971, 986–89 (2000); Lisa Heinzerling, *Regulatory Costs of Mythic Proportions*, 107 *YALE L.J.* 981, 984–85 (1998); Cass R. Sunstein, *Incommensurability and Valuation in Law*, 92 *MICH. L. REV.* 779, 784 (1994). See generally Laurence H. Tribe, *Policy Science: Analysis or Ideology?*, 2 *PHIL. & PUB. AFF.* 66 (1972).

14. There is an extensive literature criticizing the use of economic efficiency as a normative criterion. See, e.g., Adler & Posner 1999, *supra* note 4, at 191 (“[B]ecause Kaldor-Hicks [efficiency] is,

The internal critiques, in contrast, sidestep these incommensurability and normative criterion problems by accepting as a given premise the desirability of assessing policies primarily, or even solely, through aggregating the affected persons' willingness to pay to enjoy or to avoid their consequences. These critics instead address the cost-benefit methodology on its own economic efficiency-oriented terms. The critics point out a number of valuation problems which, alone or in combination, may render futile attempts to quantify costs and benefits in an objective manner because they simply do not reflect the analyst's personal ethical and political preferences.¹⁵ The valuation problems noted by these critiques include the difficulty in determining appropriate discount rates for use in evaluating future consequences,¹⁶ the problem of determining whether offer prices or asking prices are the appropriate measure of willingness to pay in instances where they diverge in magnitude,¹⁷ and the dependence of the magnitude of both offer and asking prices on the existing distribution of wealth.¹⁸ Those internal critiques commonly emphasize the severe data availability limitations facing cost-benefit analysts, including the difficulty of often having inadequate data to confidently establish affected persons' willingness to pay for even the known impacts of a policy,¹⁹ as well as the common problem of scientific uncertainty as to both the scope and magnitude of the likely consequences of many policies or programs.²⁰

I have previously attempted to contribute to both branches of this critical literature.²¹ In this Article, however, I will not address any of these external or internal critiques that I have noted. I will instead elaborate upon a different internal critique of the cost-benefit methodology, one which has been underappre-

taken as a moral principle, unsound, [cost-benefit analysis] cannot be justified by reference to Kaldor-Hicks."); Ackerman & Heinzerling, *supra* note 9, at 1567-68; Gregory Scott Crespi, *The Mid-Life Crisis of the Law and Economics Movement: Confronting the Problems of Nonfalsifiability and Normative Bias*, 67 NOTRE DAME L. REV. 231, 234-37 (1991) [hereinafter Crespi 1991]; MARK KELMAN, A GUIDE TO CRITICAL LEGAL STUDIES 114-50 (Harvard Univ. Press 1987). See generally Symposium, *Efficiency as a Legal Concern*, 8 HOFSTRA L. REV. 485 (1980) (collecting several articles on economic efficiency).

15. See Adler & Posner 1999, *supra* note 4, at 202-03. There is an extensive literature advancing such internal critiques of the cost-benefit methodology. See, e.g., Sinden, *supra* note 7, at 205-10 (citing many articles presenting such critiques). See generally Richard L. Revesz, *Environmental Regulation, Cost-Benefit Analysis, and the Discounting of Human Lives*, 99 COLUM. L. REV. 941 (1999) (criticizing commonly used approaches to discounting future loss of life consequences).

16. There is extensive literature relating to the question of the choice of discount rates by which to discount future impacts in cost-benefit analysis. See generally Revesz, *supra* note 15.

17. See Gregory Scott Crespi, *Valuation in Cost-Benefit Analysis: Choosing Between Offer and Asking Prices as the Appropriate Measure of Willingness to Pay*, 39 J. MARSHALL L. REV. 429, 432 (2006) [hereinafter Crespi 2006]. See generally Russell Korobkin, Note, *Policymaking and the Offer/Asking Problem Price Gap: Towards a Theory of Efficient Entitlement Allocation*, 46 STAN. L. REV. 663 (1994).

18. Sinden, *supra* note 7, at 206-07.

19. Sinden, *supra* note 7, at 208-10 (discussing several problems which call into question the validity and reliability of various methods of measuring willingness to pay).

20. See *id.* at 205-12.

21. See, e.g., Crespi 2006, *supra* note 17; Crespi 1991, *supra* note 14.

ciated and which has significant adverse implications for the validity of assessing of policy consequences by the willingness to pay yardstick used in cost-benefit analysis. My criticism is that conventional cost-benefit analyses do not address what I will refer to as the "endogeneity problem." That is to say, they do not confront the difficulties involved in valuing policies under those circumstances when one of their consequences is a significant alteration of the preferences of a substantial number of people, or even an alteration of the fundamental genetic identities of the members of distant future generations.

The conventional simplifying analytical assumptions made by virtually all cost-benefit analysts is that while the policy under consideration may well affect the circumstances facing persons in numerous ways, it will not affect those persons' underlying preference structures,²² nor affect the fundamental genetic identities of the members of future generations.²³ In other words, personal identities and preferences are assumed to be established exogenously and not endogenously determined by the policy under consideration. This assumption provides a stable reference point for use in valuing the consequences of the policy.

Such an assumption greatly facilitates policy analysis, but what if this is not the case? What if a particular policy significantly alters preference structures of a substantial proportion of the affected population, or fundamentally alters the genetic identity of the members of future generations, or both? The errors introduced by basing an analysis on such simplifying assumptions when they are not satisfied may well be too high a price to pay for such analytical tractability.²⁴ I will in this Article explore this possibility.

22. See Samuel Bowles, *Endogenous Preferences: The Cultural Consequences of Markets and other Economic Institutions*, 36 J. ECON. LIT. 75, 75 (1998) ("Markets and other economic institutions . . . influence the evolution of values, tastes and personalities. Economists have long assumed otherwise; the axiom of exogenous preferences is as old as liberal political philosophy itself."); Albert O. Hirschman, *Against Parsimony: Three Easy Ways of Complicating Some Categories of Economic Discourse*, 1 ECON. & PHIL. 7, 10 (1985) ("Economists often . . . think of citizens as consumers with unchanging or arbitrarily changing tastes in matters of civics as well as commodity-oriented behavior. This view tends to neglect the possibility that people are capable of changing their values."); George J. Stigler & Gary S. Becker, *De Gustibus Non Est Disputandum*, 67 AM. ECON. REV. 76, 76-77 (1977) ("[O]ne may usefully treat tastes as stable over time [N]o other approach of remotely comparable generality and power is available."); Kenneth Dau-Schmidt, *Legal Prohibitions as More Than Prices: The Economic Analysis of Preference Shaping Policies in the Law*, in LAW AND ECONOMICS: NEW AND CRITICAL PERSPECTIVES 153 (Robin Paul Malloy & Christopher Braun eds., Peter Lang Publ'g 1995) [hereinafter Dau-Schmidt 1995] ("The traditional economic analysis of law assumes that people's preferences are exogenous.").

23. I am not aware of any cost-benefit analyses that have even attempted to address this problem.

24. See Bowles, *supra* note 22, at 75 ("[By making the assumption of exogenous preferences] the scope of economic inquiry is thereby truncated in ways which restrict its explanatory power, policy relevance, and ethical coherence. If preferences are affected by the policies or institutional arrangements we study, we can neither accurately predict nor coherently evaluate the likely consequences of new policies or institutions without taking account of preference endogeneity."); Kenneth Dau-Schmidt, *An Economic Analysis of the Criminal Law as a Preference-Shaping Policy*, 1990 DUKE L.J. 1, 17 (1990) ("It is becoming increasingly apparent that the failure to address the malleability of preferences seriously limits the explanatory power of economic analysis.") [hereinafter Dau-Schmidt 1990].

The endogeneity of preferences and the endogeneity of identity problems are closely related, but there are nevertheless enough differences between their impacts on cost-benefit analysis that they merit separate treatment. Therefore, I will first address in some detail in Part II of this Article the substantial problems posed for cost-benefit analysis by the possible endogeneity of preferences, and I will then offer my suggestions as to how those problems might be better addressed. My overall conclusion here is that while these problems are indeed difficult, they appear to be surmountable through making appropriate adjustments in the cost-benefit valuation methodology. I will then turn in Part III to the more dramatic and far more imposing analytical problem posed for cost-benefit analysis by not merely the possibility but the virtual certainty that any social policy whatsoever will alter the fundamental genetic identity of many, and eventually all, members of future generations. The genetic identities of those future persons are in general not exogenous but are endogenously determined by the policies that we pursue, and this fact has momentous and perhaps even fatal consequences for the conduct of cost-benefit analyses. In Part IV I will present a brief overall conclusion.²⁵

I. THE ENDOGENOUS PREFERENCES PROBLEM

A. A General Introduction

The questions presented for policy valuation when preferences are endogenous with regard to the policy at issue have significant practical consequences, since it is clear that some policies will have as one of their impacts the substantial alteration of many of the affected persons' preferences.²⁶ Environmental policies are one area where this is particularly likely to be the case. Many persons may never have experienced certain environmental amenities, and because of their lack of familiarity with the character of those amenities, they may value them very modestly relative to their costs and therefore oppose policies designed to bring them into existence.²⁷ Once such environmental

25. I have separately addressed certain aspects each of these problems in some of my recent work. See Gregory Scott Crespi, *Incorporating Endogenous Preferences in Cost-Benefit Analysis*, 17 PENN. ST. ENVTL. L. REV. 157 (2009) (addressing the endogenous preferences problem in the environmental law context); Gregory Scott Crespi, *The Fatal Flaw of Cost-Benefit Analysis: The Problem of Person-Altering Consequences*, 38 ENVTL. L. REP. NEWS & ANALYSIS 10703 (2008) [hereinafter Crespi 2008] (addressing the problem of endogeneity of identity in the environmental law context); Gregory Scott Crespi, *What's Wrong with Dumping Radioactive Wastes in the Ocean? The Surprising Ethical and Policy Analysis Implications of the Problem of Person-Altering Consequences*, 37 ENVTL. L. REP. NEWS & ANALYSIS 10873 (2007) [hereinafter Crespi 2007] (addressing the ethical issues raised in the environmental law context by the problem of endogeneity of identity). I hope in this Article to extend that prior work and to demonstrate that these problems are closely related to one another as separate aspects of the more general endogeneity problem.

26. See Cass R. Sunstein, *Endogenous Preferences*, *Environmental Law*, 22 J. LEGAL STUD. 217, 220 (1993) ("[T]here is sometimes no such thing as a fully acontextual preference, and that preferences are endogenous to existing legal policy, including the setting of the legal entitlement.") [hereinafter Sunstein 1993].

27. See *id.* at 236–37 (giving an example of what Sunstein refers to as "adaptive preferences").

policies have been implemented, however, and people begin to experience those amenities, some of those people may well reassess their preferences in light of their greater awareness and understanding of the benefits provided by those amenities and may now accord them a higher value that exceeds their cost. A willingness to pay-based cost-benefit analysis valuation of such a policy might turn out very differently if the policy is valued by the yardstick of the post-policy implementation preference structures than if it is valued by the pre-policy implementation preference structures.

More generally, many policies have educational and character-shaping objectives,²⁸ which may sometimes even be the primary goal of the policy at issue.²⁹ If those policies are successful in achieving this objective, some persons who originally opposed those policies may, on the basis of their new attitudes and preferences, now regard them as being cost-justified. In these instances, the choice between using pre-policy implementation or instead post-policy implementation preference structures for valuations may prove outcome-determinative.

If a policy significantly alters the preference structures of a substantial number of persons, and if one wishes to incorporate this preference structure endogeneity into the valuation calculations, some difficult analytical questions are presented. Should the pre-policy implementation preference structures be used to value the policy's consequences, or should instead the different post-policy implementation preference structures be used to make those valuations? If the latter, what role, if any, should be played by any transitional post-policy implementation preference structures that may exist for a period of time before the final stable post-policy implementation preference structures come into being?³⁰

Under the simplifying exogenous preferences assumption, the same preference structures are assumed to exist both before and after the implementation of

28. See Dau-Schmidt 1995, *supra* note 22, at 153 ("[There are] various areas in the law . . . in which legal prohibitions are not merely intended to act as a price on the proscribed behavior, but are also intended to influence the underlying preferences of the sanctioned parties and other members of society.").

29. See Dau-Schmidt 1990, *supra* note 24, at 17 ("Preference shaping . . . is an important human endeavor. It has been identified as a primary or secondary goal of childrearing, education, religion, advertising, public service announcements, legislation, and . . . criminal punishment. Although economists might find it useful to assume that preference-shaping processes are exogenous to their analysis of traditional markets, when economists expand their analysis to social institutions that are more intimately related to the preference-shaping processes, either affecting or being affected by them, this assumption should be relaxed . . . Although it complicates empirical tests and introduces greater subjectivity via the social welfare analysis, relaxing the assumption that preferences are exogenous promises greater understanding of many social phenomena."); Hirschman, *supra* note 22, at 10 ("A principle purpose of publicly proclaimed laws and regulations is to stigmatize antisocial behavior and thereby influence citizens' values and behavior codes. This educational, value-molding function of the law is as important as its deterrent and repressive functions.").

30. In some cases temporary, transitional preference structures may be created for some persons by the policy's implementation, and these transitional preference structures may persist for some time before the final, stable, post-policy implementation preference structures come into being. If so, these transitional preference structures would be additional candidates for inclusion in a valuation framework. See, e.g., GARY S. BECKER, ACCOUNTING FOR TASTES 20 (Harvard Univ. Press 1996).

the policy. Analysts making this assumption for valuation purposes generally fail to make clear in their analyses whether those preference structures are being utilized for valuation purposes because they are the pre-policy implementation preference structures, with the fact that they are also the post-policy implementation preference structures not being germane, or whether they are being utilized because they are the post-policy implementation preference structures, with the fact that they are also the pre-policy implementation preference structures not being germane. There is no real need for them to address this question as to the underlying rationale of their valuation approach if the preference structures are assumed to be invariant.

Under endogenous preferences circumstances, however, a rationale must be advanced that would justify the choices made among these differing preference structures for valuation purposes. A modest body of literature exists that notes the significance of the endogenous preferences problem for policy valuation efforts³¹ and to a limited extent attempts to address the problem.³² However, that scholarship is incomplete and has not definitively resolved the conundrum, and as noted above, most cost-benefit analysts simply ignore the implications of the possibility of endogenous preferences in their work.³³

In this Part II of this Article, I will discuss how policy impacts should be valued in those instances when there are significantly differing pre-policy implementation and post-policy implementation preference structures, and possibly also distinct transitional preference structures, for a substantial number of affected persons. In order to more clearly focus my analysis on these specific questions, I will limit its scope in certain regards. First, I will take as a given the premise that people have at any point in time coherent preferences, and that their preference structures can be accurately ascertained and then used for determining those persons' willingness to pay for policy benefits or to avoid policy costs.³⁴ Given these assumptions, the main question at issue is how best

31. See, e.g., *id.* ("The endogeneity of preferences would appear to play havoc with traditional approaches to welfare evaluations of economic outcomes.").

32. See generally *id.*; Bowles, *supra* note 22; Dau-Schmidt 1995, *supra* note 22; Sunstein 1993, *supra* note 26.

33. See, e.g., a recent and otherwise rather comprehensive discussion of cost-benefit analysis written by John Graham, who served as the Director of the Office of Information and Regulatory Affairs within the Office of Management and Budget from 2001–2006, in which he does not even mention the problems posed for the validity of cost-benefit analysis by the possibility of endogenous preferences. John D. Graham, *Saving Lives Through Administrative Law and Economics*, 157 U. PA. L. REV. 395 (2008).

34. Not every commentator agrees with this assumption. In particular, Cass Sunstein has argued extensively that people's preferences often "do not, in any simple way, antedate the process that is used to elicit them . . . preferences and values can be a function of methods of elicitation, or construction." Cass R. Sunstein, *How Law Constructs Preferences*, 86 GEO. L.J. 2637, 2652 (1998) [hereinafter Sunstein 1998]. Under those circumstances, valuations achieved through use of those preferences would inevitably be suspect as being biased by the manner of elicitation. Sunstein has also argued that even apart from those biases introduced by measurement efforts, "sometimes there are no acontextual preferences with which to do normative or descriptive work," and under those circumstances "policy-makers cannot simply identify preferences and try to satisfy them since the preferences are influenced

to choose among these preference structures to obtain a policy valuation that most meaningfully reflects the preferences of the affected persons.

Second, I will consider only non-paternalistic approaches to resolving the endogenous preferences valuation problem. I will not explore the possibility of taking a paternalistic approach that might utilize for valuation purposes whichever of those two or more actual preference structures are regarded by the analyst as best reflecting the “true” interests of the affected persons, those interests being somehow defined apart from their actual preference structures. Nor will I consider the even more paternalistic tact of utilizing for valuations some other hypothetical “ideal” preference structure that the analyst believes that the affected persons “should have” to further their “true” interests, and that differs from all of the actually observed preference structures.³⁵

I will also not address in this Article the important but conceptually distinct question of whether offer prices or instead asking prices should be utilized to establish the magnitude of willingness to pay once the appropriate preference structures for valuation purposes have been chosen.³⁶ I will, however, address in this Article to a limited extent the nature and severity of the measurement difficulties involved in ascertaining these different preference structures, and how these difficulties may as a practical matter call for some departure from theoretically ideal valuation procedures. I will focus this particular discussion upon how those measurement difficulties may vary with regard to whether pre-policy implementation or instead post-policy implementation preference structures are being ascertained, and with regard to whether the cost-benefit analysis is being conducted on a prospective or instead a retrospective basis.

Let me here briefly summarize my overall conclusions with regard to the endogenous preferences problem. In my opinion, under those circumstances where preference structures are to a significant extent endogenous with regard to the policy at issue, that policy should as a general matter be evaluated solely with regard to the post-policy implementation preference structures, rather than by the yardstick of the pre-policy implementation preference structures. If there are any transitional preference structures they should be used for valuing those particular policy consequences that occur while those transitional preference structures are in existence. The use of an endogenous preference valuation framework in

by law, and since there is no way to identify the preferences that would exist in the absence of law.” Sunstein 1993, *supra* note 26, at 221.

35. I will instead accept as a premise the desirability of utilizing persons’ actual preference structures for valuation purposes in a non-paternalistic fashion, and will consider only the questions that arise under endogenous preference circumstances where there is more than one plausible candidate for this role. I recognize that this is a contestable approach. Some commentators have argued that there may often not be any such “actual” preference structures that are not biased by the method of elicitation that can be applied for making valuations without the need for prior and inevitably paternalistic corrections of those biases, that all revealed preference structures are contextual and influenced by law and thus any attempt to ascertain peoples’ preference structures for valuation purposes is also an inherently paternalistic enterprise. *See, e.g.*, Sunstein 1998, *supra* note 34, at 2652.

36. *See, e.g.*, Crespi 2006, *supra* note 17; *see also* Korobkin, *supra* note 17.

this manner when conducting cost-benefit analyses of policies that may alter preferences may upon occasion have significant consequences for the results of those analyses. This recommended endogenous preferences valuation framework may, however, need to be modified upon occasion to reflect preference structure estimation difficulties and to counter potential analyst bias, particularly when conducting prospective rather than retrospective cost-benefit analyses.

I will proceed as follows in the remainder of Part II of this Article. In Part II.B, I will briefly discuss the efforts of several other scholars to address the problem of incorporating endogenous preference into cost-benefit analysis valuation efforts. I will then, in Part II.C, offer my own thoughts regarding how this problem can be best addressed as a theoretical matter and what adjustments in that theoretical approach are likely to be necessary to accommodate practical measurement limitations. In Part II.D of the Article, I will then present a brief summary of my conclusions with regard to the endogenous preferences problem.

B. Prior Efforts to Address the Problem of Endogenous Preferences in Cost-Benefit Analysis

The most substantial contributions in this area have been provided by Gary Becker, Cass Sunstein, Kenneth Dau-Schmidt, and Samuel Bowles. Let me briefly discuss in turn the contributions of each of these scholars.

1. The Becker "Extended Utility Function" Approach

The problem of valuing policy consequences under circumstances of endogenous preferences can be analytically dealt with in a number of facially different but conceptually equivalent ways. For example, one could redefine any situation where a policy would alter preference structures as instead being a situation where the affected persons' preference structures were exogenously determined and *not* altered by that policy, but where there were now some additional policy impacts specified that were just significant enough in magnitude to generate the same overall policy valuation that would result under an endogenous preference assumption if the post-policy implementation preference structures were utilized there for making the valuations, but without those newly specified policy impacts being included in the valuation calculations. In other words, one could redefine the preference changes caused by a policy under an endogenous preferences assumption as instead being newly specified policy impacts of an appropriate magnitude under an exogenous preferences assumption, and still reach the same policy valuation results. This type of alternative conceptual framework for dealing with endogenous preferences has been suggested by the Nobel Prize-winning economist Gary Becker in his well-known book ACCOUNTING FOR TASTES.³⁷

37. See BECKER, *supra* note 30, at 3–23.

Becker there defines the concept of an exogenously determined “extended utility function”³⁸ that would also include two additional factors affecting a person’s utility that are qualitatively different from the availability of ordinary goods and services. He labels these new factors “personal capital” and “social capital,” respectively.³⁹ The amount of each of these capital stocks held by an individual at any point in time would then affect the total utility that the person would obtain from the availability of ordinary goods and services, and from the other circumstances facing that individual, at that point in time.⁴⁰ The consequences of a policy for each affected individual at any point in time would then be partly a function of the changes the policy has made in the amount of these personal and social capital stocks. Those individual valuations of the policy’s consequences at each point in time would then each be time-discounted to a present value as of the time of the implementation of the policy, and then these time-specific valuations would be aggregated into an overall assessment of the policy’s consequences for that individual. Those individual policy valuations would then be aggregated across the population to obtain an overall global policy assessment.⁴¹ Becker argues that the conventional exogenous preference assumption can be usefully and credibly retained in this manner for cost-benefit analytical work, despite his recognition that preferences are in fact often endogenous, if the valuations of policy impacts are based upon such extended utility functions that still posit exogenous preferences but which now also take into account the policy impacts upon these two appropriately defined capital stock factors over time in a suitable manner.⁴²

I am reluctant to criticize in any way an economist of Becker’s stature, but I do not believe that his suggested approach solves the endogenous preferences valuation problem. His attempt to retain the exogenous preference assumption while mitigating its shortcomings through the inclusion of two additional factors in the utility function specification that together substitute for recognizing preference endogeneity is indeed creative and interesting. But this approach ultimately only reframes the valuation problem rather than resolving it, and

38. *Id.* at 5.

39. *Id.* at 4 (“*Personal Capital, P*, includes the relevant past consumption and other personal experiences that affect current and future utilities. *Social Capital, S*, incorporates the influence of past actions of peers and others in an individual’s social network and control system.”).

40. *Id.* at 5.

41. *See id.* at 20 (“A welfare analysis should consider not only the initial effects on utility when . . . [the affected persons may intensely dislike the policy], nor only the ultimate effects when . . . [the affected persons may like the policy], but the discounted value of both the initial and later changes in utilities that incorporates the transition between the initial and later attitudes toward [the policy] In other words, initial preferences should have no priority over final preferences in welfare analysis when policies change preferences.”).

42. *See id.* (“These difficulties [of conducting valuations where preferences are endogenous to the policy in question], however, are intrinsic not to the endogeneity of preferences but to inadequate incorporation of this endogeneity into welfare criteria. If the relevant utility function for welfare analysis includes personal and social capital, the effect on utility of . . . public policies can be evaluated without any ambiguity.”).

moreover does so in a manner that tends to obscure rather than illuminate the nature of the choices that must be made.

Let me explain this point more fully. In the endogenous preferences framework, the central question presented is whether the pre-policy implementation preference structures should be used to value a policy's consequences, or whether instead the post-policy implementation preference structures, perhaps including any transitional preference structures, should be used for making those valuations. Under Becker's proposed extended utility function approach, the analogous question posed would be whether the utility functions with the pre-policy implementation stocks of personal and social capital, as he has defined them, should be used to make the policy valuations, or instead whether the utility functions with the post-policy implementation levels of these capital stocks, and perhaps also utility functions with transitional levels of these capital stocks, should be used for making the valuations.

The close parallels between an endogenous preference approach and Becker's extended utility function approach are fairly obvious. His method simply transforms the need to ascertain the contours of the different preference structures that are in existence at each point in time, a requirement for conducting valuations under the endogenous preferences framework to the extent that each of these preference structures will play a part in deriving the final policy valuations, into the need to measure the size of the different stocks of personal and social "capital," as idiosyncratically defined by Becker, at each point in time under an exogenous preferences assumption, to the extent that these different capital stocks at different points in time will play a part in deriving the final policy valuation.

There is nothing that would prohibit one from making this transformation of analytical categories, but it is not clear what the pay-off is for doing so. One could certainly incorporate those post-policy implementation stocks of Beckerian personal and social capital that exist at any point in time into the willingness to pay-based valuation calculations in the same manner that one would incorporate the post-policy implementation quantities of goods and services that normally affect utility. If the parameters of the extended utility function and the magnitude of the changes made in the personal and social capital stocks of various individuals at each point in time were each appropriately specified so as to yield the same valuations of policy impacts as would the endogenous preference framework were one to use post-policy implementation preference structures, then Becker's approach would lead to exactly the same policy valuation results as one would obtain under the endogenous preferences framework.

However, as noted above, this analytical tact only reframes rather than resolves the central valuation questions posed by the existence of endogenous preferences. The main unanswered question under the endogenous preferences formulation of the valuation problem is whether the post-policy implementation preference structures, permanent or transitional, should be privileged over the

pre-policy implementation preference structures in deriving the policy valuation, and if so to what extent. Becker's approach specifies an alternative conceptual framework of exogenous preferences combined with newly defined personal and social capital stocks that change in response to policy impacts, rather than endogenous preference changes, as establishing the parameters that need to be established to calculate the willingness to pay-based policy valuations for each relevant period of time. But his method simply transforms the central and difficult question presented by the endogenous preferences framework into the equally difficult question of whether the pre-policy implementation or instead post-policy implementation personal and capital stocks are to be privileged in conducting a policy valuation, and if so, then what weight, if any, will be given to transitional amounts of these stocks as compared to their final, stable amounts? This new question is unfortunately left unaddressed by Becker.

2. Other Commentary Regarding Incorporating Endogenous Preferences in Cost-Benefit Analysis

Let me now turn to consider the efforts of some other scholars to address the endogenous preferences valuation problem. None of the few commentators other than Gary Becker who have considered this question in any depth, specifically Cass Sunstein,⁴³ Kenneth Dau-Schmidt,⁴⁴ and Samuel Bowles,⁴⁵ offer a framework for choosing the appropriate preference structures to use to conduct a cost-benefit valuation of the impacts of a particular policy that does not require a paternalistic assessment on the part of the analyst as to which of these preference structures best reflect the affected persons' "true" preference structure, or otherwise is more deserving of respect. Let me briefly discuss each of their contributions below, and I will then draw upon their insights and the work of Becker in Part III of this Article where I will offer my own contributions.

(i) *Cass Sunstein's Views.* As I have noted, this Article proceeds upon the premise that, under endogenous preference circumstances, one can accurately measure the pre-policy implementation preference structures, post-policy implementation preference structures, and any transitional preference structures. Given this premise, I address the question of how to appropriately choose among the different policy valuations that can be derived from each of these preference structures in a non-paternalistic fashion. Cass Sunstein in his seminal 1993 article on endogenous preferences, however, calls into question this entire enterprise.⁴⁶

There, Sunstein argues that there often are no objectively existing preference structures that can be ascertained that merit such deference, and that an analyst

43. See generally Sunstein 1993, *supra* note 26.

44. See generally Dau-Schmidt 1995, *supra* note 22.

45. See generally Bowles, *supra* note 22.

46. See generally Sunstein 1993, *supra* note 26.

therefore often has no choice but to paternalistically impose what she regards as the most appropriate preference structures to use for valuation purposes. These selected preference structures may not closely resemble any of the different pre-policy implementation or post-policy implementation preference structures noted above, which he argues may not even exist in any meaningful sense.

Sunstein's core argument is that legal rules or other social policies inevitably have effects on preferences, and consequently no preference structures can be properly regarded as pristine pre-legal and pre-political expressions of people's "true" preferences.⁴⁷ In his view, pre-policy implementation preference structures often do not merit deference for use in valuing the policy's consequences because they are "adaptive preferences,"⁴⁸ that is, artifacts of and psychological accommodations to the limitations imposed by existing law and social institutions.⁴⁹ Even though these pre-policy implementation preference structures have been shaped by previously adopted laws and policies, rather than by the policy at issue that is being evaluated in cost-benefit terms in advance of its possible implementation in Sunstein's opinion, those pre-policy implementation preference structures are often so tainted as purported measures of "true" individual preferences because of the adaptive preferences effects of existing laws and other social institutions that they do not merit deference for policy valuation purposes.⁵⁰ Sunstein gives particular emphasis to arguments that debates regarding many environmental policy questions are rendered less meaningful because they take place against the background of pre-policy implementation preference structures that are distorted by this adaptive preference effect.⁵¹

Sunstein also in his 1993 article briefly calls into question the use of post-policy implementation preference structures as a basis for policy valuation, on the general grounds that "when preferences are a function of legal rules, the rules cannot be justified by reference to the preferences."⁵² There is, in his opinion, a bootstrapping problem presented by the use of post-policy implementation preference structures; policies cannot be meaningfully evaluated or justified by those very preference structures that they have helped create.⁵³

Sunstein buttresses these arguments for not utilizing either the pre-policy implementation or the post-policy implementation preference structures for

47. *Id.* at 221-35.

48. *Id.* at 236-37.

49. *Id.* at 234-35 ("The most general point is that the preference-shaping effects of legal rules cast doubt on the idea that environmental regulations should attempt to satisfy or follow some aggregation of private preferences. . . . Sometimes there is no such thing as a prelegal or prepolitical 'preference' that can be used as the basis for decision. If this is so, government is not quite faced with a choice between respecting and rejecting private preferences. This is a misdescription of the real options.").

50. *Id.* at 254 ("Private preferences are an inadequate basis for environmental policy insofar as these are adaptive to an environmentally inadequate status quo.").

51. *Id.* at 230-42.

52. Sunstein 1993, *supra* note 26, at 235.

53. *See id.* at 221, 234-35, 237.

policy valuation purposes with the related but conceptually distinct argument that any measured preference structures, even apart from their being distorted from “true” preferences by psychological adaptations to the limitations imposed by pre-policy implementation legal and political circumstances or by the effects of the policy in question, simply do not have an objective existence independent of their means of elicitation.⁵⁴ There is something analogous to the “Heisenberg uncertainty principle” at work here; any method by which an analyst attempts to measure preference structures will strongly influence the results obtained, rendering those results suspect as indicia of “true” preferences even apart from the adaptive preferences and bootstrapping effects.⁵⁵ Sunstein’s more recent work shows that he is still firmly of this opinion as to the sensitivity of measures of preference structures to the context of and methods of their elicitation.⁵⁶

Sunstein recognizes that preferences are often endogenous to a significant extent. However, as I have noted above, he implicitly rejects the core premise of this Article that it may be possible to establish an appropriate framework for selecting among the different pre-policy implementation and post-policy implementation preference structures for use in valuing a policy in a non-paternalistic fashion, because he rejects all of these preference structures as potentially being too biased and/or unreliably ascertained to use for policy valuation purposes. He therefore does not offer a proposed resolution of the endogenous preferences valuation problem in the terms that I have framed it in this Article. Some of his insights, however, are nevertheless germane to the enterprise of attempting to resolve the valuation problem viewed in these terms, and I will draw upon those insights in Part III of this Article.

Sunstein continues in his more recent work to flatly reject the idea that policy valuations can or should be carried out in a non-paternalistic manner that is grounded in expressed or otherwise revealed preferences.⁵⁷ He instead endorses the more paternalistic approach of disregarding these preferences, both pre-policy implementation preferences and post-policy implementation preferences, choosing instead to value policies with regard to their efficacy “to move people in welfare-promoting directions”⁵⁸ that are determined to be so desirable by some external criteria other than those persons’ preferences. His attempt to justify this approach in the face of the usual strong criticisms made of governmental policymakers who choose to disregard expressed or otherwise revealed

54. See Sunstein 1998, *supra* note 34, at 2652.

55. See *id.*

56. See Cass R. Sunstein & Richard H. Thaler, *Libertarian Paternalism Is Not an Oxymoron*, 70 U. CHI. L. REV. 1159, 1161 (2003) (“[I]n many domains, people lack clear, stable or well-ordered preferences. What they choose is strongly influenced by details of the context in which they make their choice, for example default rules, framing effects . . . and starting points. These contextual influences render the meaning of the term ‘preferences’ unclear.”).

57. See *id.* at 1164 (“The first misconception is that there are viable alternatives to paternalism.”).

58. *Id.* at 1162.

preferences to purportedly achieve greater social welfare⁵⁹ is to me rather unconvincing, but that is the subject of another article.

(ii) *Kenneth Dau-Schmidt's Views.* Kenneth Dau-Schmidt has written two important pieces relating to endogenous preferences, each of which focuses primarily, although not exclusively, on issues raised by the criminal law.⁶⁰ His primary emphasis in these two articles is making clear the importance and pervasiveness of the preference-shaping effects of many laws and social policies,⁶¹ and determining when preference-shaping policies are to be preferred over policies that only shape opportunities.⁶² He was not focusing upon my inquiry regarding how to properly conduct a cost-benefit assessment of a policy that has both opportunity-altering and preference-shaping consequences, a topic that he addresses only tangentially. He does, however, offer some insights that are useful for one making this effort.

Dau-Schmidt recognizes that the existence of endogenous preferences with regards to a policy undermines the application of conventional social welfare criteria to evaluate that policy,⁶³ including the willingness to pay criterion that is generally utilized in cost-benefit analysis under the exogenous preference assumption. The closest that he comes to taking an explicit position as to which of the various pre-policy implementation and post-policy implementation preference structures should be utilized for valuation purposes in cost-benefit analysis is the following, somewhat cryptic statement that he offers in his 1990 article

59. See *id.* at 1166 ("Once it is understood that . . . a form of paternalism cannot be avoided, and that the alternatives to paternalism (such as choosing options to make people worse off) are unattractive, we can abandon the less interesting question of whether to be paternalistic or not, and turn to the more constructive question of how to choose among the possible choice-influencing options. To this end we make two general suggestions. First, programs should be designed using a type of welfare analysis, one in which a serious attempt is made to measure the costs and benefits of outcomes (rather than relying on estimates of willingness to pay). Choosers should be given more choices if the welfare benefits exceed the welfare costs. Second, some results from the psychology of decisionmaking should be used to provide *ex ante* guidelines to support reasonable judgments about when consumers and workers will gain most by increasing options."). This quote clearly articulates the controversial paternalistic premise that governmental officials, supported by "reasonable guidelines" that are grounded in modern theories of psychology, can be trusted to accurately assess the "true" social costs and benefits of policies without having to rely upon the expressed willingness to pay for or to avoid policy consequences of the affected persons to make those assessments.

60. See generally Dau-Schmidt 1995, *supra* note 22; Dau-Schmidt 1990, *supra* note 24.

61. See, e.g., Dau-Schmidt 1995, *supra* note 22, at 158–68; Dau-Schmidt 1990, *supra* note 24, at 16.

62. See, e.g., Dau-Schmidt 1995, *supra* note 22, at 168–70; Dau-Schmidt 1990, *supra* note 24, at 14–24.

63. There are a variety of possible social welfare criteria that can be derived from individual utility functions that could be used for policy assessment. The approach utilized by cost-benefit analysis in aggregating willingness to pay across all affected persons is only one of the many possible options. Many of the proposed social welfare criteria, however, are grounded in some fashion upon individual preferences as to possible states of society. All social welfare functions of this type are faced with the need to deal somehow with the endogenous preferences problem in making their assessments of policies that alter preferences. See Dau-Schmidt 1990, *supra* note 24, at 6–8.

regarding what he views as the minimally acceptable properties of an adequate social welfare function:

To ensure a rational social preference ordering [in light of the 'impossibility theorem' results demonstrated by Kenneth Arrow], economists usually assume that interpersonal comparisons of the intensity of individuals' preferences are possible, and that the intensity of individuals' preferences is reflected in the social preference ordering. If one allows preferences to vary, then it is also necessary to assume that intertaste comparisons of the intensity of an individual's preferences are possible, and that intertaste intensity is also reflected in the social preference function. To account for these interpersonal and intertaste comparisons, the social welfare function has been modeled as a summation of individual utility functions which reflect the intensity of preference in their measure of utility.⁶⁴

I have assumed for the purposes of this article that the pre-policy implementation preference structures, the post-policy implementation preference structures, and any transitional preferences structures can each be accurately ascertained, and the primary question here considered is how to combine the policy valuations derived from each of these different preference structures for overall policy valuation purposes. Under this assumption, the willingness of any individual to pay for a given policy can be precisely determined under each of these preference structures, thus satisfying Dau-Schmidt's suggested minimal requirements for a valuation methodology that it be able to make interpersonal and intertaste comparisons.

These minimal requirements, however, do not mandate that any particular choices be made among the different policy valuations derived from each of these different preference structures in making the overall policy valuation. However, the major thrust of Dau-Schmidt's work is to emphasize the importance of recognizing the endogeneity of preferences in areas such as the criminal law where preference-shaping is often one of the primary goals of the policies undertaken, and to offer suggestions regarding how to assess when such preference-shaping measures are appropriate to undertake. It is consequently most unlikely that he would embrace a valuation methodology that totally ignored the post-policy implementation preference structures and any transitional preference structures, as this approach would be tantamount to making the exogenous preference assumption that he rejects as particularly inadequate in those contexts.

(iii) *Samuel Bowles' Views.* Samuel Bowles is a prominent economist who has emphasized the importance of taking into account preference endogeneity

64. *Id.* at 7-8.

for both descriptive and normative purposes.⁶⁵ While recognizing the advantages of the conventional exogenous preferences assumption for analytical tractability and for placing limits upon excessive paternalism,⁶⁶ Bowles also recognizes that under many circumstances the lack of realism of this assumption undercuts its usefulness.⁶⁷

Bowles' work parallels that of Dau-Schmidt in that he attempts to emphasize the importance of recognizing preference endogeneity in many important policy-making contexts. He also suggests the general contours of the comprehensive research program that he believes is called for to increase our currently inadequate understanding of the processes of preference formation and alteration,⁶⁸ and describes the practical significance that such greater understanding of those processes would have for economic and social policy.⁶⁹ But while Bowles generally notes that preference endogeneity has implications for policy evaluation,⁷⁰ he does not offer any specific suggestions germane to this Article's inquiry as to how the different pre-policy implementation and post-policy implementation preference structures, once determined, should be utilized for cost-benefit analysis valuation purposes. His comments do, however, suggest caution in giving substantial weight to post-policy implementation preference structures in making policy valuations under circumstances where the uncertainties regarding the accuracy of the forecasting of those future preference structures are sufficiently great to raise concerns of covert paternalistic biases being imposed by analysts.⁷¹

C. My Approach to Making Valuations in Cost-Benefit Analysis Under Endogenous Preferences Circumstances

Before I offer my own thoughts for the reader's consideration as to the most appropriate theoretical framework for the incorporation of endogenous prefer-

65. See Bowles, *supra* note 22, at 75 ("If preferences are affected by the policies or institutional arrangements we study, we can neither accurately predict nor coherently evaluate the likely consequences of new policies or institutions without taking account of preference endogeneity.")

66. See *id.* at 102 ("The implicit premise that policies . . . do not affect preferences has much to recommend it: the premise provides a common if minimal analytical framework applicable to a wide range of issues of public concern, it expresses a prudent antipathy toward paternalistic attempts at social engineering of the psyche. It modestly acknowledges how little we know about the effects of economic structure and policy on preferences, and it erects a barrier both to ad hoc explanation and to the utopian thinking of those who invoke the mutability of human dispositions in order to sidestep difficult questions of scarcity and social choice.")

67. See *id.* at 105 ("Realism, however, cannot be among the virtues invoked on behalf of the exogenous preferences assumption A broader concept of market failure is thus required, one encompassing the effects of economic policies and institutions on preferences and for this reason more adequate for the consideration of an appropriate mix of markets, communities, families, and states in economic governance.")

68. See *id.* at 102-03.

69. See *id.* at 104-05.

70. See *id.* at 104 ("[P]reference endogeneity gives rise to a kind of market failure and suggests a reconsideration of some aspects of normative economics.")

71. See *id.* at 102.

ences into cost-benefit analysis, I would like to first briefly point out a few practical considerations that may have some significance for making proper use of the results of the theoretical inquiry. I will then present my theoretical analysis and conclusions, and later revisit these practical considerations and discuss the extent to which they may impose limitations upon the use of my suggested valuation methodology.

1. Practical Considerations in Applying an Endogenous Preferences Valuation Framework

(i) *Under Some Circumstances the Choice Among Alternative Preference Structures Is Mooted.* First, there are some instances where even though a policy may significantly alter the preference structures of a substantial number of persons, it will not be necessary for a cost-benefit analyst to choose between valuations derived from the pre-policy implementation preference structure and those derived from the post-policy implementation preference structure. This will be the case whenever the net benefits calculations under either set of preference structures lead to the same conclusion as to the merits of the policy at issue relative to those of the other policy alternatives. In other words, if the valuation calculations are invariant in the sense that they would lead to endorsement (or rejection) of the policy in question relative to the alternatives when valuations are done with regard to either of these two preference structures, then the issue is rendered moot and no choice among them need be made.⁷² For example, with regard to advertising expenditures which presumably have at least some effect in altering consumer preferences as well as providing information, it has been argued that it is generally unnecessary to choose between using pre-advertising preferences or post-advertising preferences in order to value the effects of the advertising expenditures, because the results will usually not be significantly affected by the choice.⁷³

As a general matter, however, when there are significant differences in the contours of the pre-policy implementation and post-policy implementation preference structures for a substantial number of affected persons, a cost-benefit analyst might well reach different conclusions as to the relative merits of the policy in question as compared to the possible alternatives depending upon

72. However, even if all of these preference structures would yield identical policy recommendations as compared to the merits of the policy alternatives, a person following Cass Sunstein's paternalistic recommendations might reject all of those preference structures and substitute instead for valuation purposes another set of preference structures which when used for valuations in the analyst's opinion better facilitates achieving social welfare. See generally Sunstein & Thaler, *supra* note 56.

73. Avinash Dixit & Victor Norma, *Advertising and Welfare*, 9 BELL J. ECON. 1, 2 (1978) ("[I]n general, advertising that is beneficial at the margin as judged by preadvertising tastes will also be beneficial according to postadvertising tastes, and advertising that is harmful according to the latter will also be harmful when judged by the former."). This sweeping claim is certainly a counter-intuitive result, since one would expect advertisers to attempt to alter preferences in favor of the advertised goods and services rather than merely limit themselves to providing information as to their properties and availability and that such preference-shaping efforts may or may not be successful."

which set of preference structures was used for valuing the policy's consequences. Under those circumstances, a choice between these preference structures for valuation purposes will have to be made.

(ii) Preference Structure Estimation Difficulties May Limit the Ability to Implement the Theoretically Appropriate Valuation Approach. There are generally difficulties of one sort or another that will be encountered in estimating the contours of preference structures, particularly with regards to preference structures that do not yet exist, or no longer exist, at the time of estimation. In addition, there are differential difficulties encountered in estimating preference structures depending upon whether the cost-benefit analysis is being undertaken prospectively in advance of a policy's implementation or is instead being carried out retrospectively in order to assess whether a policy that has been implemented has had the desired results. Let me address each of these problems in turn.

(a) General Preference Structure Estimation Difficulties

It is of course often difficult even under the conventional simplifying exogenous preference assumption for cost-benefit analysts to ascertain the contours of the applicable preference structures so as to calculate the willingness to pay-based valuations of the consequences of the policies under consideration. There is often a lack of the revealed preference data that would be most useful for ascertaining these preferences, and the various methods used by analysts to ascertain preferences in the absence of such data are open to criticism.⁷⁴ This problem obviously becomes more severe under endogenous preferences assumptions where two or more different preference structures may need to be ascertained. The practical difficulties involved in determining the contours of each of perhaps several preference structures suggest that it may often be infeasible to attempt to ascertain the contours of transitional (and perhaps short-lived) preference structures, and that the most that can reasonably be accomplished in many instances would be to estimate only two distinct preference structures—the pre-policy implementation preference structures and the final stable long-term post-policy implementation preference structures—and then choose which of these two preference structures should be utilized to value the policy's consequences. It may even in some instances be impossible to estimate with any real confidence the post-policy implementation preference structures, forcing the analyst to in effect make his valuations under an implicit exogenous preference assumption, even when that assumption is recognized to be false.

74. See, e.g., SUNSTEIN 1993, *supra* note 26, at 218–20.

(b) Prospective and Retrospective Cost-Benefit Analyses Present Different Kinds of Problems for Determining the Applicable Preference Structures

Another point worth keeping in mind is that the practical difficulties presented by the need to determine the contours of the relevant preference structures will be different both in kind and in severity depending on whether the cost-benefit analysis is being done prospectively or instead retrospectively. Cost-benefit analyses are most commonly done prospectively as part of the process of reaching a decision as to whether to implement a particular policy. Under those circumstances, the pre-policy implementation preference structures already exist and can ideally be empirically ascertained.⁷⁵ However, the subsequently resulting post-policy implementation preferences structures do not yet exist before the implementation of the policy. They will have to be projected solely on the basis of estimates of the impacts that the policy will have on the pre-policy implementation preference structures. Given our limited understanding of the processes of preference formation and alteration,⁷⁶ such projections will be relatively uncertain and subject to additional sources of error that go beyond the difficulties inherent in ascertaining existing preference structures.

This need to forecast the contours of post-policy implementation preference structures raises a real concern that an analyst might be tempted to utilize result-oriented projections of preference structure alterations that are formulated with the hidden agenda of either favoring or disfavoring the policy at issue.⁷⁷ Given the inherent uncertainty of such projections of future preference structures, these kinds of analyst manipulations may be difficult or impossible to recognize and discredit prior to the implementation of the policy, by which time it may be too late to effectively revisit the decision.

This practical concern as to the potential for analyst bias that would be difficult or impossible to police adequately might call for giving relatively greater weight to valuations derived from the more easily ascertained existing pre-policy implementation preference structures, and correspondingly giving relatively lesser weight to the more subjective and uncertain projections of post-policy implementation preferences structures, than would be justified by purely theoretical considerations were there not such significant differentials in the measurement difficulties presented by the need to establish the contours of the alternative preference structures. In a case where the estimates of post-policy implementation preference structures are particularly uncertain, this concern might even justify relying solely upon the pre-policy implementation preference structures – in effect reimposing the assumption of exogeneity of preferences for the policy at issue under circumstances where this assumption is

75. At least, in theory, this can be done, although determining the relevant individual preference structures in the absence of revealed preference data can be very difficult.

76. See Bowles, *supra* note 22, at 102–03.

77. See *id.*

recognized to be false—in order to avoid the possibility of being misled by analyst bias.

Cost-benefit analyses are, however, sometimes done retrospectively rather than prospectively in an attempt to determine whether a policy has been successful in achieving its objectives, though this effort of revisiting policies with the advantage of hindsight is not undertaken nearly as often as some commentators would prefer.⁷⁸ Under those retrospective analysis circumstances, the post-policy implementation preference structures exist at the time at which the analysis is conducted and can ideally be empirically ascertained. In addition, the pre-policy implementation preference structures and any transitional preference structures, while they would no longer exist at the time of the analysis, would presumably have left some historical record and can be more accurately ascertained than can post-policy implementation preference structures in the prospective analysis context where they can only be forecast.

This difference in analytical context suggests that there should be somewhat less concern about the possibility of hidden analyst bias in preference structures estimation for retrospective cost-benefit analyses than for prospective analyses. In other words, the differential difficulties presented with regard to establishing the contours of preference structures that do not exist at the time that the analysis is conducted are more serious in the prospective cost-benefit analysis context than they are in the retrospective analysis context. Therefore, the practical concerns about avoiding potential sources of analyst bias that might call for giving greater weight to the valuations derived from the pre-policy implementation preference structures than to those valuations that are derived from estimated post-policy implementation preference structures in the prospective analysis context are more significant than are the comparable concerns that would call for giving greater weight to valuations derived from post-policy implementation preference structures than to those derived from the pre-policy implementation preference structures in the retrospective analysis context. It may, therefore, be the case that the policy valuations derived from endogenous preferences assumptions should, under some circumstances, perhaps be calculated somewhat differently for prospective cost-benefit analyses than for retrospective analyses, even absent any theoretical reasons for such disparate treatment.

2. A Proposed Methodology for Combining Valuations Obtained Through the Use of Different Preference Structures

To provide a context for my subsequent analysis, I ask the reader to assume for the sake of argument that a cost-benefit analyst wishes to evaluate the impacts of a proposed policy, either prospectively or retrospectively, under circumstances where the policy will have significant effects upon the preferences of a substantial number of people. In other words, assume that there

78. See, e.g., Hahn & Tetlock, *supra* note 7, at 77–78.

initially exists one set of preference structures that characterizes those persons' pre-policy implementation preferences, but that another and different stable set of preference structures will characterize those persons' preferences after the policy has been implemented for some period of time, and perhaps one or even a series of transitional preference structures will also characterize those persons' preferences at various times during the transition period to their final stable post-policy implementation preference structures. Assume also for now that each of these preference structures can be accurately ascertained so that the uncertainty of estimation and analyst bias issues discussed above do not arise. Finally, further assume that one is taking a non-paternalistic stance under which one is committed to deferring to the extent possible to the preferences of the affected persons. Given these assumptions, what combination of policy valuations derived from each of these preference structures would be the most appropriate theoretical framework to use for valuing the policy impacts in accordance with the underlying willingness to pay principle that essentially defines cost-benefit analysis?

I have concluded that at least in most instances the post-policy implementation preference structures should be used to value policy consequences, rather than the pre-policy implementation preference structures, for both prospective and retrospective cost-benefit analyses. The most intuitively plausible justification for this choice is the obvious fact that the post-policy implementation preference structures embody the actual preferences that will exist when the policy consequences are experienced by the affected persons. This justification is given further support by Dau-Schmidt's argument that to use the pre-policy implementation preference structures to value those policies that are intended partially or even primarily to alter preferences would be to completely miss the point of those policies.

Sunstein, as discussed above, is for various reasons critical of using either pre-policy implementation or post-policy implementation preference structures to value policies, but his arguments certainly cannot be read to favor the use of the pre-policy implementation preference structures over the post-policy implementation preference structures for making valuations; if anything, they would suggest the contrary. Neither Becker nor Bowles specifically address this choice among preference structures for valuation purposes question in their work, except that Becker briefly notes that "initial preferences should have no priority over final preferences in welfare analysis when policies change preferences,"⁷⁹ and that Bowles notes the potential for analyst bias in the estimation of post-policy implementation preference structures, an issue I will later address when considering practical limitations upon the use of an endogenous preferences valuation methodology.

I offer the "at least in most instances" qualification in my recommendation above because one can certainly conceive of extreme circumstances where

79. BECKER, *supra* note 30, at 20.

valuations that are made solely on the basis of post-policy implementation preference structures can be criticized as a means of attempting to justify undesirable policies on the basis of distorted preferences that the policies themselves have created.⁸⁰ As a practical matter, however, most policies do not appear to directly distort preferences but merely provide more information that allows people the option to alter their preferences, but does not force them to do so. I therefore do not think the bootstrapping concern adverted to by Sunstein significantly undercuts the argument for relying solely upon the post-policy implementation preference structures for policy valuation in the usual cases.

This approach of using post-policy implementation preference structures to value policies that alter preferences will tend to favor those environmental policies and other preference-altering policies for which the post-policy implementation preference structures are more favorable to the policy in question when weighing the trade-offs between policy benefits and costs. In some instances, this effect could be outcome-determinative; a policy option that is regarded as inferior when assessed by the yardstick of the pre-policy implementation preference structures may prove to be the preferred alternative when assessed with regard to the post-policy implementation preference structures. I have argued in this Article that under these circumstances the latter, more favorable assessment of the policy in question is the correct cost-benefit assessment.

My decision to endorse the use of the post-policy implementation preference structures for valuation purposes raises the derivative question as to what weight to give to transitional preference structures, where they exist, relative to the final stable post-policy implementation preference structures. My conclusion here is in general accord with Becker's discussion of the proper roles that time-discounting and preference structure duration should play in policy valuation and should be relatively uncontroversial. Policy consequences that occur while transitional preference structures exist should be valued on the basis of those transitional preference structures and then time-discounted to a present value based on when those consequences occur. Once the final stable post-policy implementation preference structures have come into existence, all subsequent policy consequences should be valued in accordance with those preference structures and then time-discounted to a present value as appropriate.

Under this discounting approach, the longer a transitional set of preference structures exists, other things being equal, the more weight they will be given in the overall policy valuation. This seems most sensible. In addition, the time-discounting applied to the valuations derived from each of the applicable post-policy implementation preference structures to convert them all to present

80. One can imagine extreme examples of this situation. Would anyone accept as a justification for imposing forced lobotomies on people the fact that post-lobotomy, those persons seemed content with their new lot in life? As a less extreme example, a policy that facilitated access to addictive drugs might well be favorably regarded by subsequent addicts who had initially opposed those policies pre-addiction. Should the valuations by the addicts then be regarded as definitive as to the value of the addictive drug access policy?

value terms will render those valuations commensurate for the purpose of aggregation into an overall policy valuation.

This endogenous preferences valuation framework that I recommend here on theoretical grounds may, however, need to be practically modified in some instances to reflect the preference structure estimation difficulties and potential analyst biases that I have discussed above. I will address these issues in the next subsection of this Article.

3. Recognizing Practical Limitations on an Endogenous Preferences Valuation Methodology

As noted above, there may sometimes be severe difficulties involved in estimating the contours of transitional preference structures. In such instances, it may be necessary as a practical matter to make the simplifying assumption that the final stable post-policy implementation preference structures will immediately come into existence upon the implementation of the policy. In those instances where the transitional preference structures are short-lived or where they do not differ significantly from the final stable post-policy implementation preference structures, or both, this simplifying assumption is unlikely to introduce significant error into the valuation calculations.

Most cost-benefit analyses are done prospectively, and in some instances the analyst may unfortunately not only be unable to determine the contours of any transitional preference structures, but also unable to estimate with any confidence even the contours of the final stable post-policy implementation preference structures. Under these circumstances, there is no ideal way to proceed, and a difficult choice is presented as to the best analytical approach.

One could direct the analyst to proceed with the policy valuation exercise on the basis of the pre-policy implementation preference structures, assuming that they can be empirically ascertained with reasonable accuracy, essentially utilizing the exogenous preferences framework in a situation where this assumption is known or at least believed to be inaccurate. On the other hand, one could direct the analyst to persevere with the endogenous preferences valuation approach in the face of this difficulty and use the best estimate of the post-policy implementation preference structure(s) that she can ascertain. This latter tact of course raises the possibility that the analyst may incorporate into the cost-benefit analysis biased estimates of these preference structures that are designed to justify a particular result that she favors, and that will be difficult or impossible for others to recognize as such and criticize. This problem of potential and effectively unreviewable analyst bias will, as previously discussed, be much less of a concern in the retrospective cost-benefit analysis context where the post-policy implementation preference structures exist and can presumably be empirically estimated with some reasonable accuracy, and where there is no need under my suggested valuation approach to undertake the more difficult exercise of estimating the contours of the former pre-policy implementation preference structures on the basis of historical evidence.

D. Conclusions Regarding Incorporating Endogeneity of Preferences

Cost-benefit analysis is the most important policy evaluation technique now used in American public sector decision making. It is therefore crucial that these analyses be properly conducted. The conventional assumption made by cost-benefit analysts, usually implicitly rather than explicitly,⁸¹ is that individual preference structures are not altered by any of the policies that are under consideration. This simplifying exogenous preferences assumption is not always satisfied, however, and in some instances the preference structures of a substantial proportion of the people who are impacted by a policy are endogenous in that they are also significantly altered by that policy. Under those endogenous preferences circumstances, an important question is presented as to whether valuations of the impacts of the policies called for by the cost-benefit methodology based on willingness to pay, should be calculated with regard to the pre-policy implementation preference structures, or instead with regard to the different post-policy implementation preference structures, including any transitional preference structures that may exist for a period of time, or perhaps with regard to some combination of the above.

Several prominent scholars have previously addressed aspects of this inquiry, including most importantly Gary Becker, Cass Sunstein, Kenneth Dau-Schmidt, and Samuel Bowles. While each of those writers has made a substantial contribution to the effort, the endogenous preferences valuation question has not yet been definitively resolved and virtually all cost-benefit analysts continue to ignore the implications of the possibility of endogenous preferences in their work.⁸² In this Article I have assessed the merits of Becker's extended utility function valuation approach, and of the various valuation suggestions offered by Sunstein, Dau-Schmidt, and Bowles. Within a framework of hopefully reasonable assumptions that I have made as to the ability of analysts to ascertain the contours of preference structures, I have offered my own thoughts regarding how the endogenous preferences valuation problem can be best addressed in a non-paternalistic fashion.

My main conclusion is that the willingness to pay-based valuations of policies provided by cost-benefit analyses should be derived solely from the post-policy implementation preference structures, as a general matter. If there are any transitional preferences structures, they should also be used to value those policy consequences that occur while those transitional preference structures are in existence. This valuation framework recognizes the important fact that post-policy implementation preferences are the actual preferences that will exist when the policy consequences are experienced by the affected persons, recognizes that many policies will alter and are even designed to alter preferences, and gives proper weight to the timing of policy consequences and the duration of various post-policy implementation preference structures.

81. See, e.g., Graham, *supra* note 33.

82. *Id.*

This recommended cost-benefit valuation framework will be more favorable to those environmental and other policies that have as one of their consequences the alteration of preferences in a manner that favors the policy at issue, and in my opinion that favoritism is justified for the reasons discussed above. As I have also discussed, however, this theoretical endogenous preferences valuation framework may need to be modified at times to accommodate preference structure estimation difficulties and to counter potential analyst bias, particularly with regard to prospective cost-benefit analyses. My hope is that cost-benefit analysts will come to recognize the shortcomings of invariably assuming that preferences are exogenous with regard to the policies under consideration, and will give more thought to ways of incorporating preference endogeneity into their analyses.

II. THE ENDOGENEITY OF IDENTITY PROBLEM

A. A General Introduction

While the possible endogeneity of preferences presents significant problems for cost-benefit analysis, endogeneity of identity poses a far more serious and perhaps even fatal challenge to that methodology. The nature and severity of this problem is even less widely appreciated by analysts than is the endogeneity of preferences problem.

In my previous work I have referred to the fact that the fundamental genetic identities of the members of future generations are impacted by the current social policies that we pursue—in other words, that their fundamental identities are not exogenous to those policies, but are endogenously determined by them—as the “problem of person-altering consequences.”⁸³ I will continue to use that phrase to describe the endogeneity of identity problem in this Article. I will discuss the nature of these person-altering consequences in detail below. For now let me merely state succinctly that what I am referring to by this phrase is the fact that the outcome of a successful act of human reproduction is radically contingent. Which particular one of the hundreds of millions of sperm that are released in an ejaculation will unite with the female egg, if any, is a very uncertain event. Even the slightest difference in the timing or any other aspects of a reproductively successful act of intercourse will almost certainly lead to a different sperm-egg fusion, and therefore ultimately to the birth of a genetically different individual than would otherwise have been born. The person now conceived and born will be a different individual in the most fundamental genetic sense.

The consequences of this simple fact are momentous. Any social policy that is significant enough in its direct or indirect impact on human behavior to lead to even a single different sperm-egg fusion taking place will create a genetically different individual than the person that would have been born absent the implementation of the policy. Even the most minor and local policy will surely

83. Crespi 2008, *supra* note 25; Crespi 2007, *supra* note 25.

have that much impact on someone's behavior. Over time, as that genetically different individual is born and matures, and over their life influences numerous other people in major or minor ways, this will result in an exponentially spreading cascade of fundamental genetic changes in the population of individuals subsequently conceived. After a relatively short transitional period, in a historical sense, the identity of all individuals that are conceived and born over the rest of eternity will be fundamentally different from who they would have been in the absence of the policy.

In other words, one dramatic impact of any significant policy measure will be the elimination of all members of the population of distant future generations that would have been conceived and born absent the policy, and their replacement by an entirely different group of people. Yet another way to describe the situation in more technical economic language is to say that we face a pervasive endogeneity of identity problem. The identity of future persons is not determined exogenously, but is instead determined after a transitional period of time endogenous to the policies pursued. These person-altering consequences would, from the perspective of the affected individuals, completely dwarf in significance the combined effect of all other policy consequences.

In conventional cost-benefit analyses these important endogenous person-altering consequences are invariably overlooked.⁸⁴ The typical cost-benefit analyst calculates both the benefits and the costs of the policy at issue by the willingness to pay of the affected persons, as compared to the reference point of a hypothetical, counterfactual baseline scenario of a world in which the exact same persons exist but without the policy's impacts. This procedure is tantamount to an implicit assumption that personal identity is exogenous; that the same future population of individuals will exist whether or not the policy is implemented. Such an assumption is not merely implausible but is demonstrably false, and is equivalent to simply ignoring those person-altering consequences. The calculation of costs and benefits relative to an unachievable and thus meaningless baseline reference scenario renders the conclusions of such an analysis essentially irrelevant to the real choices at hand among the actual consequences that are possible to achieve through the alternative policies under consideration.

It might at first appear that this analytical problem could be solved simply by more realistically specifying the hypothetical baseline scenario used as a reference point for valuing the impacts of the policy at issue. This baseline scenario could be specified in a manner that recognizes that different future persons would exist were the policy to be implemented, and generate its pervasive and eventually universal person-altering consequences, than were the policy not to be implemented. The valuations then would be done in a more accurate fashion that took into account these endogenous person-altering consequences that would ensue if the policy were implemented. Unfortunately, further reflection suggests that the problem is not so easily fixable.

84. See, e.g., Graham, *supra* note 33.

The basic analytical conundrum presented is that if one attempts to so incorporate these person-altering consequences into a cost-benefit analysis, rather than simply ignoring them, the valuation calculations become so unwieldy and imprecise as to essentially be indeterminate. As I will discuss below, it then becomes difficult (if not impossible) to avoid the conclusion that *any* policy option that is pursued will generate massive (if not infinite) aggregate future net benefits of very uncertain magnitude, relative to the appropriately specified baseline scenario, even if these future benefits are time-discounted at relatively high discount rates. Such a sweeping conclusion that massive future benefits of highly uncertain magnitude will likely result no matter which course of action is pursued seems intuitively implausible, to say the least. It certainly does not provide meaningful guidance to policymakers for discriminating among alternative courses of action.

Once a cost-benefit analyst eats of the apple of the tree of knowledge, i.e., recognizes that all policies have pervasive and eventually universal and eternal person-altering consequences, she is put into in a real bind with no good choices available within this methodological framework. Continuing to ignore those consequences, given their overwhelming significance to the persons affected, would disregard the willingness to pay valuation principle that underlies the cost-benefit approach, and is no longer an option if one wants to reach results that are relevant to the actual choices at hand. However, once one incorporates those person-altering consequences into the analysis, there does not appear to be any way to meaningfully estimate in a willingness to pay-based manner and then compare the size of the massive benefits to future generations that would result under each of the various policy options so as to provide useful guidance.

The problem is fundamental since cost-benefit analysis is simply the systematic application to decision making of the basic willingness to pay valuation principle of welfare economics, which is in turn logically derived by straightforward reasoning from conventional and widely embraced secular, consequentialist ethical premises.⁸⁵ In light of the seemingly insurmountable problems that the willingness to pay-based valuation framework faces in meaningfully assessing the significance of person-altering consequences, it may simply be the case that cost-benefit analysis should no

85. By the phrase "secular premises" I refer to ethical premises that are derived from reflections on the human condition that are agnostic with regard to the existence of a supreme supernatural being of ethical relevance. I will not address in this Article the difficult question as to whether there is a sufficient basis in one or more of the mainstream religious traditions for recognizing an ethical obligation to the members of distant future generations. By the phrase "consequentialist premises" I refer to the ethical premise that actions have ethical relevance only to the extent that they have consequences for the rights or interests of specific persons, and that actions have no ethical relevance in and of themselves apart from those consequences. I will consider briefly in Part II.C. of this Article whether there exist any sufficient secular but non-consequentialist grounds for asserting that we have any ethical obligations to the members of distant future generations, or to the human race as a whole, that exist apart from any ethical obligations grounded in the consequences of our conduct for specific persons. However, my arguments in this Article are addressed primarily to the relatively narrow yet practically quite important question as to whether and if so how cost-benefit analyses that are based upon conventional secular and consequentialist ethical premises should be conducted if one wishes to incorporate person-altering consequences into those assessments.

longer be regarded as a useful analytical tool.⁸⁶ If one cannot reach meaningful conclusions and policy recommendations solely on the basis of conventional ethical premises and their willingness to pay-based valuation corollary, then analytical efforts may henceforth have to incorporate a broader range of non-consequentialist or even theistic normative criteria as alternative or supplementary bases for offering recommendations to policymakers.⁸⁷ This is a rather troubling thought, since once one goes beyond giving weight only to the policy impacts upon specific individuals as measured by their willingness to pay any consensus as to the appropriate evaluative criteria to apply will be most difficult to achieve.

In the remainder of Part III of this Article I will proceed as follows. In Part III.B below I will explain in more detail what I mean by the phrase “person-altering consequences.” In Part III.C I will discuss the serious problems encountered in attempting to incorporate person-altering consequences into cost-benefit analyses in the conventional willingness to pay manner. I will then discuss a couple of alternative ways that one might attempt to incorporate those consequences into the cost-benefit valuations in a meaningful fashion that still respects the willingness to pay principle, and will point out the serious shortcomings of each of these efforts that renders them inadequate responses to the problem. I will then offer a few preliminary thoughts regarding what new normative criteria might be utilized to assess policy impacts upon future generations, and how the conclusions reached through application of those criteria might be combined with the willingness to pay-based valuations of policy impacts upon existing persons. Part III.D will present a brief and rather pessimistic overall conclusion regarding the viability of the cost-benefit methodology in light of this problem even if it is so augmented with additional normative criteria.

B. *The Problem of Person-Altering Consequences*

The noted British philosopher Derek Parfit first articulated in 1976⁸⁸ a simple yet profound insight that philosophers have since labeled “the Non-Identity

86. Such analyses may, however, continue to have considerable utility as persuasive rhetorical devices when addressing audiences that are not sophisticated enough to recognize how much the validity of the cost-benefit methodology is undercut by the pervasiveness of person-altering consequences, whether or not those consequences are taken into account.

87. I have addressed more fully the ethical implications of the problem of person-altering consequences, and the application of this line of thought to major environmental policy decisions such as radioactive waste disposal approaches and global warming mitigation efforts, in Crespi 2007, *supra* note 25; see also Crespi 2008, *supra* note 25. I also address in Crespi 2008 to a minor extent the cost-benefit analysis concerns presented by this problem that are discussed more fully in this Article. Crespi 2008, *supra* note 25, at 10885–86.

88. Derek Parfit, *On Doing the Best for Our Children*, in *ETHICS AND POPULATION* 100–15 (Michael Bayles, ed. 1976) [hereinafter Parfit 1976]. Gregory Kavka argues that Parfit’s insight was also discovered independently at approximately the same time by Robert Adams and Thomas Schwartz. Gregory S. Kavka, *The Paradox of Future Individuals*, 11 *PHIL. & PUB. AFF.* 93–112 (1982) (citing Robert M. Adams, *Existence, Self-Interest, and the Problem of Evil*, 13 *NOÛS* 53 (1979), and also citing Thomas Schwartz, *Obligations to Posterity*, in *OBLIGATIONS TO FUTURE GENERATIONS* 3–13 (Richard Sikora & Brian Barry, eds., 1978). Parfit further discusses his insights in Derek Parfit, *Future Generations, Further Problems*, 11 *PHIL. & PUB. AFF.* 113,

Problem,”⁸⁹ and which I will refer to in this Article as either the problem of endogeneity of identity or the problem of person-altering consequences.⁹⁰ This insight calls into serious question whether we have any ethical obligations at all to distant future generations that can be justified on the basis of conventional secular, consequentialist ethical premises.⁹¹ It also renders inadequate any analytical efforts that overlook those consequences, even though, as I will discuss below, it is difficult if not impossible to conduct meaningful cost-benefit analyses that do incorporate those consequences. While this problem has fostered substantial (though inconclusive) discussion among philosophers and other scholars over the last three decades at an abstract, academic level regarding its ethical significance,⁹² its dramatic practical implications for policymakers in general and cost-benefit analysts in particular have not yet been adequately addressed.⁹³

115–17 (1982) [hereinafter Parfit 1982], in DEREK PARFIT, *REASONS AND PERSONS* 351–80 (1984) [hereinafter PARFIT 1984], and in Derek Parfit, *Comments*, 96 *ETHICS* 832 (1986) [hereinafter Parfit 1986].

89. Parfit later labeled this problem as the “Non-Identity Problem,” PARFIT 1984, *supra* note 88, at 378, and it is generally so described by other academic philosophers. *See, e.g.*, Doran Smolkin, *Towards a Rights-Based Solution to the Non-Identity Problem*, 30 *J. SOC. PHIL.* 194, 194 (1999); David Wasserman, *The Nonidentity Problem, Disability, and the Role Morality of Prospective Parents*, 116 *ETHICS* 132 (2005). The problem also is described by some other scholars as the “Parfit Paradox.” *See, e.g.*, Kavka, *supra* note 88, at 95 (“[Parfit’s] argument poses a . . . Paradox of Future Individuals”); Edith Brown Weiss, *What Obligation Does Our Generation Owe to the Next? An Approach to Global Environmental Responsibility: Our Rights and Obligations to Future Generations for the Environment*, 84 *AM. J. INT’L. L.* 198, 204 (1990) (referring to this insight as “Derek Parfit’s famous paradox”); Lothar Gundling, *What Obligation Does Our Generation Owe to the Next? An Approach to Global Environmental Responsibility: Our Responsibility to Future Generations*, 84 *AM. J. INT’L. L.* 207, 210 (1990) (referring to this insight as “Parfit’s paradox”). Those scholars who regard Parfit’s insight as posing a paradox commonly state the question that he poses along the lines of “How can we owe a duty to future persons if the very act of discharging that duty wipes out the very individuals to whom we allegedly owe that duty?” *See, e.g.*, Anthony D’Amato, *What Obligation Does Our Generation Owe to the Next? An Approach to Global Environmental Responsibility: Do We Owe a Duty to Future Generations to Preserve the Global Environment?*, 84 *AM. J. INT’L. L.* 190, 191 (1990). I prefer to pose the problem as a non-paradoxical though difficult question of determining the ethical implications of policies that have among their other long-term effects pervasive person-altering consequences; the elimination of the existence of all yet-unborn future persons who would have been born absent the policy’s impacts, and the birth instead of a different set of future persons.

90. In my opinion, Parfit’s own “Non-Identity Problem” label is more apt than the “Parfit Paradox” label because the question is not really a paradox so much as it is a conceptual problem regarding the existence of ethical obligations. However, Parfit’s label obscures somewhat the precise nature of the problem for those who are not academic philosophers and are not familiar with the problem and the body of scholarship that it has engendered. I therefore will use in this Article the more straightforward descriptive phrases “endogeneity of identity” or “person-altering consequences.”

91. *See supra* note 85.

92. *See, e.g.*, Adams, *supra* note 88; Schwartz, *supra* note 88; Kavka, *supra* note 88, James Woodward, *The Non-Identity Problem*, 96 *ETHICS* 804–31 (1986); Smolkin, *supra* note 89.

93. *But see* Crespi 2007, *supra* note 25, where I have attempted to contribute to such a fuller assessment. There is a legal literature of modest size and scope that addresses some of the implications of the problem of person-altering consequences, but that literature fails to fully incorporate the insights of the philosophers who have addressed the matter. *See* Aaron-Andrew P. Bruhl, *Justice Unconceived: How Posterity Has Rights*, 14 *YALE J. L. & HUMAN.* 393, 397 (2002) (“[T]he topic of future generations’ rights has spawned a growing literature—or, rather, at least two separate literatures, one in law and the other in philosophy, with very little interaction between the two.”). Most of the relevant legal literature

In this Article I will not attempt to fully articulate or resolve the complex philosophical arguments that have been offered regarding the problem of person-altering consequences, although I will reference much of that literature for those philosophically-oriented readers who wish to later explore this problem in a more rigorous and systematic fashion.⁹⁴ I will instead discuss the problem in a more condensed and straightforward manner that is intended to be helpful to academics in other fields and practicing lawyers who are not deeply versed in these technical philosophical debates, but who nevertheless wish to better understand the nature of the problem of person-altering consequences and its implications for practical policy making, particularly with regard to its implications for the conduct and relevance of cost-benefit analyses.

Parfit has clearly been the primary instigator of and contributor to discussions of the difficulties involved in assessing the person-altering consequences of policies through several works that he published over the 1976–1986 decade.⁹⁵ The most significant of these efforts were his seminal 1976 article⁹⁶ and his

focuses on the specific person-altering consequences issues raised by assisted reproductive technologies; whether persons born with birth defects as a result of such technologies but who would not have otherwise been born have standing to claim that they were thereby injured by a “wrongful life” tort. *See, e.g.*, Carter J. Dillard, *Rethinking the Procreative Right*, 10 YALE HUM. RTS. & DEV. L. J. 1 (2007); John A. Robertson, *Procreative Liberty and Harm to Offspring in Assisted Reproduction*, 30 AM. J. L. & MED. 7 (2004); Eric Rakowski, *Who Should Pay for Bad Genes?*, 90 CAL. L. REV. 1345 (2002); Phillip G. Peters, *Harming Future Persons: Obligations to the Children of Reproductive Technology*, 8 S. CAL. INTERDISC. L. J. 375 (1999); Michael Laudor, *In Defense of Wrongful Life: Bringing Political Theory to the Defense of a Tort*, 62 FORDHAM L. REV. 1675 (1994). There are only a few articles in the legal literature that consider the significance of person-altering consequences in other contexts. *See, e.g.*, Amy J. Sepinwall, *Responsibility for Historical Injustices: Reconceiving the Case for Reparations*, 22 J. L. & POLITICS 183 (2006) (considering the significance of person-altering consequences for the validity of the claims made by the descendants of slaves for reparations payments); Lukas H. Meyer, *The Palestinian Refugees and the Right of Return: Theoretical Perspectives: Historical Injustice and the Right of Return*, 5 THEORETICAL INQ. L. 305 (2004) (considering the significance of person-altering consequences for the validity of the claims made by the descendants of displaced Palestinian refugees for a right to return to their ancestral homeland). I am not aware of any prior attempts to more broadly assess the significance of person-altering consequences for the conduct and relevance of cost-benefit analysis in all contexts whatsoever, although Douglas Kysar has recently recognized that those consequences do pose “deep conceptual challenges” to any analytical method such as cost-benefit analysis “that is framed in terms of the rights, preferences, or interests of particular individuals.” Kysar, *supra* note 8, at 37.

Two other prominent law professors who have addressed the problem of person-altering consequences at least tangentially in their work are Daniel Farber and Jeffrey Gaba. Daniel A. Farber, *From Here to Eternity: Environmental Law and Future Generations*, 2003 ILL. L. REV. 289, 302 n.69 (2003) (citing quotes made by Derek Parfit on this topic); DANIEL A. FARBER, *ECO-PRAGMATISM: MAKING SENSIBLE ENVIRONMENTAL DECISIONS IN AN UNCERTAIN WORLD* 151 n.38 (1999) at (same); Jeffrey M. Gaba, *Environmental Ethics and Our Moral Relationship to Future Generations: Future Rights and Present Virtue*, 24 COLUM. J. ENVTL. L. 249, 258 n.24 (1999).

94. See Crespi 2007, *supra* note 25, for a more comprehensive review of this literature.

95. Parfit 1976, *supra* note 88; Parfit 1982, *supra* note 88; PARFIT 1984, *supra* note 88; Parfit 1986, *supra* note 88.

96. Parfit 1976, *supra* note 88. In this 1976 essay Parfit uses the hypothetical situation of a woman deciding whether to postpone becoming pregnant until she recovers from an illness that would result in any child conceived being born with a handicap to illustrate the person-altering consequences of a decision for persons who as a result will now *not* be born. *Id.* at 100–01. Parfit notes that if the

more comprehensive 1984 book *REASONS AND PERSONS*,⁹⁷ but he has also made other contributions to this debate.⁹⁸ Parfit's seminal insight is that virtually any

pregnancy is postponed and her child is later conceived after the illness is cured, the child that would initially have been born will not now be born without the handicap, but instead will never be born at all. *Id.* at 101. It is instead someone else with a different genetic endowment that will be born without the handicap.

Whether a handicapped child is better off for not being born is an impossible question to answer in abstract, general terms. It seems plausible that most if not all handicapped persons would prefer their lives to nonexistence, particularly if the handicaps are of lesser severity. A mother who decides to become pregnant while suffering such an illness and who therefore bears a handicapped child cannot be criticized on the basis of the consequences for the child unless we assume that the child would have preferred nonexistence to being born. *Id.* Such a sweeping and counterintuitive assumption about the nature of the preferences of future persons is unwarranted. Similarly, it is far more plausible that future people would prefer existence even with severe environmental constraints to nonexistence.

To illustrate the effects of policies with person-altering consequences on people who *will* be born as a result Parfit in this Article poses the hypothetical situation of a policy measure that would have only positive effects upon existing persons, but that would also have very adverse effects for future persons. *Id.* at 101–02. He notes that absent the implementation of the policy those particular future persons would never have been born, and argues that they would prefer living subject to the adverse effects of the policy at issue to the alternative of never having been born. *Id.*

97. PARFIT 1984, *supra* note 88. Parfit once again revisited the questions posed by person-altering consequences in this comprehensive 1984 book, now for the first time labeling the issue the “Non-Identity Problem,” *id.* at 378 (“This problem arises because, in the different outcomes, different people would exist. I therefore call this the Non-Identity Problem.”), and devoting an entire chapter to its analysis, *id.* at 351–379, that drew heavily upon his earlier 1982 article, PARFIT 1984, *supra* note 88. After an extended analysis of numerous hypothetical situations Parfit concludes by reasserting his position that a policy that has person-altering consequences will not be worse for those persons thereby born as a result of the policy than would be their nonexistence should the policy not have been implemented. *Id.* at 378.

Parfit takes in this book what he labels the “No Difference View;” the claim that although a policy will have person-altering consequences and therefore will not be worse for any specific individual it still could be judged to be undesirable on moral grounds. Person-altering consequences that make a policy worse for no one ultimately make no difference in a moral evaluation of that policy. *Id.* at 366–71. He also concludes his chapter on the Non-Identity Problem by reaffirming his earlier broad, aspirational claim first made in his 1982 Article, Parfit 1982, *supra* note 88, at 169–72, that it may be possible to formulate a valuation approach that appropriately addresses the problem of person-altering consequences, and which can justify moral condemnation even of policies that hurt no one. PARFIT 1984, *supra* note 88 at 377–79. He generically labels this approach “Theory X,” *id.* at 378, and states that he will later in the book attempt to formulate such a theory. *Id.* at 379 (“In what follows I will try to find Theory X.”). He predicts once again that this criterion will *not* be based upon an assessment of whether its consequences are good or bad for affected future persons. *Id.* His final and more pessimistic conclusion at the end of this book, however, is that he has again failed to formulate an approach that adequately addresses the problem of moral evaluation in the context of person-altering consequences without creating other difficulties that render the approach unacceptable, although he still optimistically believes that it might yet be possible to do so. “Though I failed to discover X, I believe that, if they tried, others could succeed.” *Id.* at 443.

Parfit notes several properties that his envisioned “Theory X” would have to satisfy to be an adequate ethical approach: “Theory X must solve the Non-Identity Problem, avoid the Repugnant and Absurd Conclusions, and solve the Mere Addition Paradox. I failed to find a theory that can meet these four requirements.” *Id.* at 443. “Most of us would believe that the Repugnant and Absurd Conclusions are what I have called them. Until we know how to avoid both conclusions, and how to solve both the Non-Identity Problem and the Mere Addition Paradox, we will have beliefs that we cannot justify, and that we know to be inconsistent.” *Id.* at 452.

98. Parfit 1982, *supra* note 88; Parfit 1986, *supra* note 88. Parfit later revisited the questions posed by the person-altering consequences of policies that he had first raised in his original 1976 article in a

human action, however slight its impacts, is likely to have at least minor effects on the timing of or other circumstances surrounding some acts of sexual reproduction, leading to different sperm-egg fertilizations than would otherwise have taken place, and consequently will over time lead to exponentially cascading consequences of a person-altering nature as now genetically different individuals mature and reproduce and in other ways influence the sexual behavior of a broader and broader circle of people.⁹⁹ After probably no more than a few decades at the most this will lead to the entire human population for all eternity now being composed of individuals that each have significantly different genetic endowments from those persons that would have existed absent that initial policy action.¹⁰⁰ The policy will thus have changed the core identity of all of those future persons; they will be different people in the most fundamental genetic sense.

later 1982 article, Parfit 1982, *supra* note 88, that was written as a companion piece to a shorter article on the topic by Gregory Kavka, Kavka, *supra* note 88, that Parfit's 1976 essay had helped to inspire. Kavka, *supra* note 88, at 93. After a long and detailed analysis of the problem, and of Kavka's modified Kantian categorical imperative-type proposals for addressing it, Parfit once again concluded that policies with person-altering consequences simply cannot be properly evaluated on the basis of whether the results of those policies are better or worse for the rights or interests of future persons. Parfit 1982, *supra* note 88 at 171–72. He again argued that "some new principle of beneficence" not yet identified that is not based on those person-affecting considerations will be needed to judge the merits of such policies. *Id.* He closed by stating that while it would be quite difficult to formulate such a new principle, it would be hasty to conclude that it was impossible to do so, because "non-religious moral philosophy is a very young subject." *Id.* at 172.

Parfit's 1986 ETHICS article, Parfit 1986, *supra* note 88, was included in an ETHICS symposium issue focusing on his 1984 book REASONS AND PERSONS, and which also included contributions by Brian Barry, Susan Wolf, Bart Schultz, Shelly Kagan, Bart Gruzalski, Arthur Kuflik and James Woodward. Parfit 1986, *supra* note 88. In that article Parfit responded in detail to each of the other symposium contributors' comments on his 1984 book. In particular, Parfit responded in some detail to James Woodward's article, Woodward, *supra* note 92, at 804–31 that specifically focused on the Non-Identity Problem. Parfit 1986, *supra* note 88 at 854–62. He there reaffirmed his longstanding "No Difference View" conclusion, most fully articulated in PARFIT 1984, *supra* note 88, at 366–71, that despite the fact that policies with adverse long-term consequences will not make things worse for any particular future individuals because of those policies' person-altering consequences, there are still moral reasons for not choosing such policies, Parfit 1986, *supra* note 88 at 854, and that these moral reasons are just as strong in spite of the fact that no particular individuals are harmed by such policies, *id.* at 855–56. Parfit once again conceded, however, that he was unable to formulate the needed "new theory about beneficence" that would justify this conclusion. *Id.* at 854.

99. See D'Amato, *supra*, note 88, at 190–92, for fuller elaboration of this point.

100. *Id.* at 191. How rapidly the person-altering consequences of a policy will proliferate, and how quickly the genetic divergence will be large enough to be of major significance to the personal identities of the persons affected, will differ from policy to policy. The identity of future persons will begin to be altered approximately nine months after the implementation of a policy once persons conceived after the policy's implementation begin to be born (I am classifying those persons conceived before a policy is implemented but perhaps affected in uterus by its consequences as "existing persons" with regard to that policy). One would expect that given the obvious sensitivity of the forming of a particular sperm-egg fusion to a great multitude of circumstances that the scope of the person-altering consequences of a policy would expand with exponential rapidity once some genetically altered individuals are born, and that even relatively isolated human communities would be impacted and then completely transformed within a few years or at most a few decades. In addition, while the initial genetic alterations resulting may in some instances be relatively minor in impact (eye color, "junk" DNA changes, etc.), arguably leaving unaffected the "identity" of the persons genetically altered in such minor fashion, the number of genetic alterations per person will also exponentially increase over time through the same

Personal identity is thus in all instances an endogenous variable. Put another way, one rather dramatic consequence of any policy measure, even one of rather limited and localized initial impact, is that over the longer-term it will eliminate the coming into existence of many and eventually all future individuals¹⁰¹ who would otherwise have been conceived and born. It will result instead in the conception and birth of an increasingly and eventually entirely genetically different group of people,¹⁰² with their genetic endowments also increasingly diverging from those of the persons who would otherwise have been born.¹⁰³ Thus, the policy will fundamentally alter the personal identities of all members of distant future generations; one vast group of potential persons will now never be conceived and will be replaced by an entirely different group of individuals. From the perspective of those affected persons, both those who will now be conceived and born as a consequence of the policy, and those who as another consequence of the policy will now never come into existence, there could not be a more dramatic impact. Therefore, these consequences—as well as the other consequences of a policy—need to be taken into account in any comprehensive assessment of its merits, whether that assessment is done through the cost-benefit methodology or otherwise.

As a matter of scientific fact,¹⁰⁴ Parfit's insight is correct and exemplifies what is commonly referred to as the "butterfly effect" of chaos theory,¹⁰⁵ where small perturbations in initial conditions can lead to massive overall systemic effects. Parfit was primarily concerned in his philosophical work on person-altering consequences with assessing their ethical implications, which he understandably found to be quite disturbing.¹⁰⁶ In this Article, however, I will address

cumulative feedback process, rapidly leading to unarguable fundamental alterations in identity of all future persons.

101. There will be a post-policy "transitional period" starting approximately nine months after the implementation of a policy during which some, but not all, individuals born will have their fundamental genetic identities altered by the a policy's consequences, before the person-altering consequences are of universal scope. The length of such a transitional period, and the proportion of individuals born with fundamentally altered identities at each point in time during this transitional period, will presumably vary from policy to policy.

102. See D'Amato, *supra* note 89, at 191.

103. Not only will the proportion of births that are policy-altered increase over time after the implementation of the policy, but the cumulative magnitude of the genetic alterations for typical individuals will also increase over time as the policy's person-altering consequences become more widely manifested and reinforce one another.

104. See D'Amato, *supra* note 89, at 192. This conclusion assumes, of course, that a person's identity is determined by their genetic endowment and/or by the physical and cultural circumstances of their lives, rather than determined by some kind of ethereal Cartesian ego or "soul" that is wholly independent of genetic characteristics or physical or cultural influences. I will assume for the purposes of this Article that if the genetic endowment of a person is significantly altered as a consequence of a policy this can be regarded as a change in that person's fundamental identity, whereas any consequence of a policy that does not significantly alter a person's genetic endowment, no matter how significant that consequence otherwise is to that person's life, does not change the fundamental identity of that person.

105. See, e.g., *id.* at 190–92.

106. Parfit himself is obviously most uncomfortable with the unavoidable implication of his insight that current policies that favor existing persons but that have adverse or even disastrous impacts upon

those ethical implications only tangentially, and will focus instead upon the implications of person-altering consequences for the conduct and relevance of cost-benefit analysis of policy options.

C. *The Implications of Person-Altering Consequences for Cost-Benefit Analysis*

1. Applying Standard Cost-Benefit Valuation Techniques to Assess Person-Altering Consequences

It is immediately apparent that to evaluate the merits of a policy that will have person-altering consequences—which I would argue includes any policy whatsoever given the inevitability of exponentially spreading genetic consequences from even initially very minor effects—through a cost-benefit analysis, it will not be adequate merely to apply a time discount to the policy's future impacts, as is now done under the conventional implicit assumption embodied by the usual specification of the hypothetical baseline scenario that the policy will not alter the preferences nor the identities of future individuals, but will only affect their wealth in some fashion. Person-altering consequences will have to be explicitly incorporated in some manner. The justifications generally offered for discounting future impacts at an appropriate discount rate are not affected by inclusion of these consequences,¹⁰⁷ but it is now also necessary to sharply differentiate between the policy's future impacts upon existing persons¹⁰⁸ over the rest of their post-policy implementation lives, which do not

future persons would nevertheless be regarded as beneficial by those future persons relative to their alternative of nonexistence if the policy is not pursued, and thus those policies cannot be criticized on the usual person-affecting basis that they would injure particular people.

[T]he long-term effects of social policies, even if clearly disastrous—even if it clearly affects for the worse—won't be worse for particular people. They are thus ignored by our principle. . . . We might claim that we should grant *less* weight to the further future. . . . But a 'person-affecting' principle gives to the further future *no* weight.

This seems indefensible.

Parfit 1976, *supra* note 88, at 102. Parfit thus demonstrates that he understands the serious problem posed by person-altering consequences for any utilitarian criterion or related measure such as the Kaldor-Hicks wealth-maximization criterion that attempts to aggregate in some fashion the impacts of policies upon the affected persons. "Such difficulties [posed by person-altering consequences] may seem to face only utilitarians. This is not so. They face most of those who give any weight to a utilitarian principle." *Id.* at 100. He is unfortunately somewhat opaque in this brief 1976 essay regarding how this problem should be resolved. He rejects the alternative of simply ignoring the exponentially cascading person-altering consequences that will generally occur when a policy is implemented, particularly given that the total number of future persons that would be born will also likely be affected as well as their individual identities. *Id.* at 103. He does state that the problem of person-altering consequences implies that the long-term consequences of policies should *not* be determined by their impacts upon the rights and interests of the affected future persons, *id.* at 102, but he does not offer an alternative valuation method.

107. I will not in this Article address the question of whether discounting the future impacts of a policy is appropriate, or if so how the appropriate discount rate should be ascertained. For discussion of these issues, see generally Revesz, *supra* note 15.

108. By the term "existing persons" I mean to refer broadly to not only those persons who are already born at the time of the policy in question's implementation, but also those persons already

include person-altering consequences, and the policy's impacts upon future persons,¹⁰⁹ which will include those consequences.

One must initially recognize that there are two distinct groups of future persons that will be affected in fundamentally different ways by the person-altering consequences of a policy. There is, first of all, the very large group of future persons who will be conceived and born over the subsequent course of history as a result of those consequences. For them, the implementation of the policy is a necessary condition of their existence. Second, there is the vastly larger group of what I will here refer to as "potential but now never to be conceived future persons" who would have been conceived and born as a consequence of our pursuing one or another of the potentially unlimited number of alternative courses of action other than the policy at issue, including the null option of taking no action, but who will not be conceived if the policy at issue is implemented.¹¹⁰

It is obvious that the hypothetical preferences of this second group of untold trillions of potential but now never to be conceived future persons should not be given any weight in a cost-benefit analysis of the policy at issue. With their very existence at stake, each of these future persons would likely regard any specific policy—other than the single policy that would result in their coming into existence—as imposing immense costs upon them,¹¹¹ resulting in a very large if not infinite aggregate cost measure for any specific policy whatsoever that would dominate any measure of benefits that is utilized. This absurd result of massive rejection of any course of action whatsoever (including the null option of taking no action) indicates that it would be a category mistake to accord standing to potential but now never to be conceived future persons in an analysis of the consequences of a policy that necessarily precludes the existence of those persons. The hypothetical preferences of all of those future persons who have the potential to exist under one policy alternative or another, but whose existence would be precluded by the specific policy measure under

conceived but not yet born at that time, as well as those members of the "transitional generations" following the policy who were conceived after the implementation of the policy and whose genetic identity has not yet been significantly altered by the spreading person-altering consequences of the policy.

109. By the term "future persons" I mean to refer broadly to those persons conceived after the implementation of the policy in question whose genetic identity has been significantly altered by the spreading person-altering consequences of the policy. For any policy there will be a transitional period of some length during which some but not all persons born will have had their genetic identities fundamentally altered by the person-altering consequences of the policy, before those consequences become universal in scope. *See also supra* note 107.

110. It is perhaps a misnomer to refer to these wholly imaginary "beings" that never will come into existence at all as being "persons" in any sense whatsoever, but for lack of a better descriptive phrase I will refer to them as such.

111. This is under the assumption that offer prices, a more conservative and constrained measure of willingness to pay, are utilized as the approach for measuring willingness to pay. These aggregate costs would likely be infinite if asking price measures rather than offer price measures were utilized.

consideration, should be ignored in assessing that policy's effects.¹¹²

But what about the first group of future persons who *will* be conceived and born post-policy, for whom the policy's implementation is a necessary condition of their existence? One would expect that at least the overwhelming majority if not all of these future persons who would owe their very existence to the implementation of a policy would, if given the opportunity, assign very high offer prices¹¹³ to the policy even were that policy to have some adverse or even catastrophic consequences for their well-being. Given that these person-altering consequences will persist for all eternity, it is of course not knowable in advance how many future persons from each era would exist and be voting in such a hypothetical referendum, let alone what the wealth endowment and precise preference structure of each of these future persons that would constrain the magnitude of their offer prices would be. It therefore will not be possible to ascertain the distribution of the costs and benefits of a policy between existing persons and future persons. What is clear, however, is that any policy with person-altering consequences, no matter how broadly catastrophic its long-term impacts, would result in truly massive benefits for those future persons¹¹⁴ who otherwise would not have been born that would completely dominate the magnitude of any adverse impacts upon existing persons,¹¹⁵ for the obvious

112. Jeffrey Gaba has insightfully likened this situation to the science fiction motif of an infinite number of universes being generated each instant as our present decisions create multiple alternative futures. Gaba, *supra* note 93, at 258 n.24. He also draws the analogy to multi-universe interpretations of the probabilistic results of quantum physics. *Id.* He concludes as do I that the adverse impacts of our policies upon this multitude of potential but now never to be conceived persons should be ignored, though not for the reason that I give that their inclusion in the analysis would lead to absurd results, but instead because their competing interests should be regarded as "cancelling out;" cost-benefit analysis should in effect be "renormalized" to eliminate such infinite values in a manner that parallels what physicists do in their quantum mechanics equations. *Id.*

Steven Landsburg in his recent and entertaining book, *MORE SEX IS SAFER SEX* (2007), also addresses to a modest extent the question of how to value the consequences of our actions for yet-unconceived future persons. *Id.* at 238–43. Landsburg recognizes that our policy choices raise moral questions with regard to their impact upon yet-unconceived future generations ("Do we have any moral obligation to account for the interests of trillions of *potential* people, who will never have the opportunity to live unless we conceive them?"), and that these questions are of practical significance for real-world policymaking. *Id.* at 243. He also recognizes the perhaps insurmountable difficulty of these questions. *Id.* at 239 ("Perhaps [we should just admit] . . . that we're incapable of being logically rigorous about issues involving the unconceived."). His analysis, however, appears to regard unconceived future persons as comprising a single large group who can either be conceived or not, depending on what course of action we pursue, rather than recognizing that they actually constitute a vast multiplicity of alternative groups of persons extending through time. A policy action leading to the conception of one group would necessarily preclude the conception of all of the other groups, necessitating the development of a framework for addressing these stark intra-group conflicts of interest were any rights for unconceived persons to be recognized. Landsburg does not address this difficulty, and consequently does not appear to understand the full significance of person-altering consequences for policy analysis.

113. It is also likely that infinite asking prices would be assigned if this is the willingness to pay measure utilized.

114. This is true even using restrictive offer price measures rather than uncapped asking price measures of these benefits.

115. I am assuming that future persons are psychologically similar to existing persons in this regard.

reason that all of the untold trillions of future persons whose hypothetical preferences are being considered would owe their very existence to the implementation of that policy. This presents a real problem for cost-benefit analysis because any policy measure whatsoever, including the null option, will now result in massive benefits of highly uncertain magnitude for the combined group of existing persons and future persons when person-altering consequences are included. Of what use, if any, would such analyses be for policymakers in choosing among alternatives?

Consider, for example, a radically present-oriented proposal to put all of our high-level radioactive wastes into ordinary steel barrels that will not provide effective long-term containment and then dump them all overboard into the Pacific Ocean.¹¹⁶ This policy would free billions of dollars of resources now devoted each year to radioactive waste storage efforts to be diverted to other pressing social needs. While those future persons born several centuries from now and thereafter may well suffer very significant adverse environmental consequences from such an action, the multi-billion dollar resource reallocations that such a policy would allow would have cascading person-altering consequences that would surely be universal in scope well before those barrels began to leak their poisons.

As previously discussed, those potential but now never to be conceived future persons who as a result of those resource reallocations will not be born should not be accorded standing in a cost-benefit analysis of the ocean waste dumping policy.¹¹⁷ The future persons that will be born as a consequence of that ocean waste dumping policy would owe their very existence to it. If they could be asked for their opinions about the policy, if they are at all like existing persons in their psychological make-up they would surely overwhelmingly (if not unanimously) prefer coming into existence, even if their lives involved grappling with a serious radioactive waste problem, to nonexistence. They would of course much prefer existence without the radioactive waste problem, were that an option that could be chosen, but the central insight of the problem of person-altering consequences is that this is not logically possible. The only choice that those future persons would be hypothetically presented with is the bundled Hobson's Choice of life with the radioactive waste problem or nonexistence, and if they are at all like existing persons they would assign very large benefits to the policy however those benefits are assessed.

The ocean waste dumping policy will therefore be very favorably evaluated by a cost-benefit analysis that incorporates person-altering consequences, since the existing persons who will surely be dead long before the wastes leak into the environment and cause adverse biological consequences would be, on balance, net beneficiaries of the large resource reallocations thereby made possible by

116. This particular hypothetical is analyzed in some detail in my earlier articles on the subject. See Crespi 2008, *supra* note 25; Crespi 2007, *supra* note 25.

117. See *supra* note 113 and the accompanying text.

the policy,¹¹⁸ and the long parade of generations of future persons for whom the policy is a necessary condition of their existence could be plausibly assumed to chime in with declarations of truly massive benefits, though of very uncertain magnitude in the aggregate. So, this ocean waste dumping policy will receive a ringing endorsement from the cost-benefit analysis. As I have already made clear, however, so would any other policy proposal, even those that are broadly disfavored by existing persons, since the cost-benefit calculations would invariably be dominated by the vast horde of members of distant future generations for whom the particular policy at issue would be a necessary condition of their existence.

The severe valuation problem posed for cost-benefit analysis by person-altering consequences is thus squarely posed. For those untold trillions of future persons whose identity will be affected by those consequences of a policy, the policy is a necessary condition of their existence. Its impacts will thus be valued very highly by those persons as against their actual alternative of nonexistence. The conventional practice of valuing the consequences of a policy as compared to the hypothetical baseline scenario of a world in which those same persons would exist, but without experiencing the policy's impacts, makes no sense at all since such an alternative scenario could not possibly occur. Willingness to pay-based assessments of benefits for future persons that are derived in such a fashion are completely arbitrary. Moreover, assessments so derived are not only arbitrary but are also biased downwards, in some instances dramatically so. This is because for some policies (such as the ocean waste dumping hypothetical that I have discussed above) many future persons may strongly prefer the unattainable scenario in which they are presumed to still exist, but without experiencing the impacts of the policy at issue, as compared to the world that would actually result for them from the policy's consequences. Under this comparative framework, those future persons would then likely assign net costs rather than very large benefits to the policy's consequences, leading in the aggregate to a perhaps massive undervaluation of the future effects of the policy as compared to its valuation if those future persons were to assess it as against their actual alternative of nonexistence. This approach would thus give far too much weight to the consequences of that policy for existing persons relative to its actual massively beneficial impacts on future persons. The current conventional approach of utilizing hypothetical baseline scenarios that ignore person-altering consequences is fatally flawed; that much is clear.

In my opinion, the failure of cost-benefit analysts to incorporate person-

118. I concede that there may well be existing persons who empathize sufficiently with the environmental problems that the ocean waste dumping policy may cause for the members of distant future generations that they would regard the policy as imposing net costs on themselves, despite the more immediate and tangible benefits that may accrue to them from the resource allocation savings. However, I feel confident that on balance the net costs this policy would impose on these unusually empathetic persons will be substantially outweighed by the net benefits for the large number of existing persons whose empathetic time horizons do not span as far into the future as several centuries or more.

altering consequences in their analyses has been primarily due to their overlooking those consequences rather than deliberately choosing to ignore them. There is apparently a general lack of familiarity among these analysts with the work of Derek Parfit and other philosophers who have wrestled with the problem of giving proper weight to these kinds of consequences. An argument can be made that the use of a demonstrably unattainable hypothetical baseline reference scenario that ignores person-altering consequences can be justified on the basis of parallels between this endogeneity of identity situation and the lesser problem presented to cost-benefit analysis by endogenous preferences. I believe, however, that this argument is very strained and unconvincing.

Let me explain how this analogy fails. One may attempt to characterize the endogeneity of identity problem presented by person-altering consequences as simply being an extreme extension of the endogenous preferences situation discussed above in Part II, different only in magnitude and not in its essential character. One might then seek guidance for evaluating policies with person-altering consequences from the ideas that have been proposed for addressing the lesser difficulties posed by endogenous preferences. However, in my opinion, this approach does not prove to be fruitful.

In the paradigmatic endogenous preferences situation, a group of persons whose circumstances have been impacted by a policy have also had their preference structures altered by the policy, although their fundamental genetic identities are assumed to be unchanged. The question posed for cost-benefit analysis, as discussed above in Part II of this Article, is whether those persons' pre-policy implementation preferences or post-policy implementation preferences should be utilized to generate their willingness to pay-based cost and benefit assessments of the policy's impacts, as compared to the reference point of the policy not being implemented and those persons' initial circumstances and preference structures both remaining unchanged. One could argue the endogeneity of identity problem may be viewed in some regards as simply an extreme extension of the endogenous preferences problem, and thus amenable to some of the same valuation adjustments discussed above.

However, this analogy between the endogenous preferences and endogeneity of identity situations is rather superficial and breaks down quickly under closer inspection. In the paradigmatic endogenous preference situation only currently existing persons are considered, and a choice need only be made as to which of the two preference structures that exist at different points in time for those persons better reflects their "true" preferences. In the endogeneity of identity situation, in sharp contrast, the future persons involved do not exist prior to the implementation of the policy and have only one preference structure, their post-policy implementation preferences, that very strongly favor the policy at issue. The use of the conventional hypothetical baseline scenario assumption, therefore, does not merely substitute an earlier and different structure of preferences held by those persons for use in obtaining those persons' valuation of the policy. That conventional assumption in this instance instead substitutes as the

reference point for comparison a hypothetical factual circumstance—those persons' existence but without the policy's consequences—that cannot possibly occur and bears no relationship to those persons' preferences, earlier or later, "true" or otherwise. The arguments made by Cass Sunstein that may justify use of pre-policy implementation preferences to value a policy in the endogenous preferences context, under some circumstances, obviously do not justify the use of a demonstrably false "the same persons will exist either way" standard of comparison in the endogeneity of identity context.

Person-altering consequences can no longer be simply ignored. But as noted, the other horn of the dilemma is that attempting to value those person-altering consequences in the usual willingness to pay-based manner unfortunately leads to the cost-benefit analysis methodology "blowing up." All policy options will now generate massive future benefits¹¹⁹ that will dominate the policies' possibly adverse effects upon existing persons, but those future net benefits are simply not measurable with sufficient precision to allow the alternative policy options to be meaningfully compared and ranked.

Is there a viable middle ground here? Can we develop a willingness to pay-based method of including person-altering consequences in cost-benefit analysis that would lead to intuitively reasonable results and meaningful discrimination among policy alternatives? Or do we face a fundamental and insoluble problem with the cost-benefit methodology in that person-altering consequences cannot credibly be ignored any longer, but also cannot be meaningfully valued in the willingness to pay-based fashion that essentially defines the cost-benefit methodology, which may require us to discard this approach and take an entirely new analytical tact?

This endogeneity of identity problem may well be fatal to cost-benefit analysis. It may simply be that because of the pervasiveness and significance of person-altering consequences, meaningful policy recommendations cannot be formulated solely on the basis of conventional secular, consequentialist ethical premises and their willingness to pay-based valuation corollary. We may have little choice but to develop new assessment methodologies grounded at least partly upon non-consequentialist normative criteria to adequately handle those consequences. However, before endorsing such a radical change, let me first explore a couple of alternative willingness to pay-based valuation approaches to including person-altering consequences in cost-benefit analysis, which one may consider in an attempt to rescue the methodology from this dilemma.

2. Consideration of Two Alternative Cost-Benefit Approaches for Valuing Person-Altering Consequences

Let me set forth and discuss two alternative valuation approaches that would incorporate person-altering consequences and each retain the conventional will-

119. In fact, they will clearly each result in infinite net benefits if asking price measures of willingness to pay are utilized.

ingness to pay valuation premise that is the very heart of the cost-benefit methodology. One approach would be to first value all of the impacts of a policy on the members of distant future generations—for whom it is a necessary condition of their existence—by their estimated offer prices, rather than by their estimated asking prices.¹²⁰ This approach would initially generate a finite—although still very large—aggregate benefit measure, rather than an analytically intractable infinite benefit measure. Then, discount those future benefits at a very high discount rate—high enough so they have an aggregate present value of essentially zero. Under this approach, the massive benefits of a policy for the members of distant future generations—for whom it is a necessary condition of their existence—would then be reduced to insignificance and thus not overwhelm the effects of the policy on existing persons, which would then essentially become the sole consideration in reaching a conclusion.

Let me consider this approach in more detail. First, the use of offer price rather than asking price measures of impacts has some plausibility. Offer prices, rather than asking prices, are now conventionally used in cost-benefit analyses to measure both the costs and benefits of the policy under consideration.¹²¹ While some observers regard this choice of valuation measures as arbitrary and imposing a sometimes severe downward bias on the numbers thereby obtained,¹²² a more in-depth analysis of the question suggests that the use of offer prices rather than asking prices can perhaps be adequately justified, at least for benefit measures, although the question of how the cost impacts of a policy are most appropriately measured is a much more difficult question whose proper resolution is still uncertain.¹²³

The use of a very high discount rate to minimize the policy benefits for future persons is, however, far more problematic. Most arguments offered for discounting future impacts are based upon either observed rates of investor time preference or social rates of return on invested capital and cannot plausibly justify the use of annual discount rates higher than, at most, 10–15%.¹²⁴ While discount rates of such magnitude would reduce to relative insignificance even very large benefits that occurred a century or more in the future,¹²⁵ they would not prevent the person-altering consequences for the many millions of people

120. Whether offer prices or instead asking prices should be used to measure costs and benefits is a very controversial question that I have elsewhere explored at length, see generally Crespi 2006, *supra* note 17, and plausible arguments can be made for the use of offer price measures of benefits. *Id.* at 464–65.

121. *Id.* at 436.

122. *Id.* at 445.

123. The issue of whether offer prices or instead asking prices are the most appropriate way to measure costs and benefits is exhaustively addressed in Crespi, *id.*, and in Korobkin, *supra* note 17.

124. I will not address in this Article the current debate regarding the appropriate choice of discount rates with regard to conventional cost-benefit analyses that ignore person-altering consequences. See generally Revesz, *supra* note 15.

125. For example, benefits that occurred 100 years in the future would by the use of a 10% annual discount rate be reduced by a factor of approximately 13,740, and benefits occurring 200 years in the future would be reduced by a factor of approximately 1,888,000!

likely to be affected within a few years (or at most a few decades) after a policy's implementation—given the great sensitivity of particular sperm-egg unions to even very minor changes in the circumstances of an act of sexual intercourse—from completely dominating the calculations.¹²⁶ Annual discount rates in the high triple digits would probably be necessary to reduce those nearer-term large future benefits to insignificance.¹²⁷

High triple-digit annual discount rates, however, cannot be grounded on any plausible theory of investor time preferences or social rates of return on capital investment, and the use of this approach to value person-altering consequences is obviously a contrived means of nominally addressing person-altering consequences in a willingness to pay-based valuation framework, while in substance ignoring those consequences so as to avoid the paralyzing computational difficulties. Moreover, this approach is perhaps even less adequate than the current conventional approach in handling the problem of person-altering consequences. The current conventional approach, as discussed above, rather than attempting to value person-altering consequences implicitly substitutes valuation of a policy's effects relative to a demonstrably unattainable baseline scenario that ignores those consequences. The suggested offer prices/very high discount rate alternative, in contrast, candidly recognizes the existence of person-altering consequences as a formal matter but then proceeds to value them at essentially zero through the high discount rate ploy. In some instances, this zero measure of impacts upon future persons could be even less accurate than the arbitrary measure that is obtained through the conventional approach, under which it is quite possible to obtain a positive rather than zero aggregate valuation of a policy's impacts on future persons relative to the hypothetical baseline scenario. Thus, it may lead, in some cases, to an even more pronounced bias towards radically present-oriented policies than the current approach.

Despite these shortcomings, one could argue that an approach that at least in principle recognizes the existence of person-altering consequences, even though it then mathematically manipulates them out of the analysis, is at least a small step in the right direction. Those important consequences are now not completely ignored at the outset, and the core willingness to pay-based valuation framework has nominally been retained. However, this approach ultimately avoids confronting the difficult valuation question in a meaningful fashion and

126. For example, benefits that occurred five years in the future would, by the use of even a 15% annual discount rate, only be discounted by a factor of approximately two. Even those benefits that did not occur for 30 years, by which time the person-altering consequences of almost any policy are likely to be universal, would only be discounted by a factor of approximately 66, leaving them likely still several orders of magnitude greater than the policy's impacts upon existing persons.

127. For example, use of a 100% annual discount rate would discount benefits that occurred five years in the future by a factor of only 32, which would still result in the very large benefits resulting from person-altering consequences dominating the analysis. However, an annual discount rate of, say, 700%, would lead to the discounting of such fifth-year benefits by a factor of 16,807, which might suffice to reduce them to an aggregate level that is insignificant relative to the policy's impacts on existing persons.

may suffer from an even stronger bias in favor of radically present-oriented policies than does the conventional approach. I conclude that this offer prices/very high discount rate alternative approach is an inadequate means of modifying cost-benefit analysis to incorporate person-altering consequences.

The second alternative valuation approach that I suggest for consideration would accomplish essentially the same result, but in a more candid manner. Under this approach, one would simply ignore all the impacts of a policy on the members of distant future generations and consider only the benefits and costs for existing persons.¹²⁸

The basic argument for taking this approach is as follows. Those members of distant future generations who would be conceived and born as a result of any particular policy choice, and who understood the narrow range of logically possible outcomes—the limited bundled choices that we face because of the policy's person-altering consequences for future generations—would ascribe very large benefits to that policy since it provides the necessary conditions for their existence. But, there is no feasible way to quantify and compare the huge aggregate benefits that would result for each of the different groups of future persons that would be conceived and born under each of our many possible policy options. Given this fact, perhaps it makes sense to simply ignore the massive beneficial impacts of each policy under consideration on the particular group of future persons that it brings into being and focus only upon the impacts upon existing persons.¹²⁹ Those massive but practically indeterminate benefits to those future persons who are born under each of the different possible policy options might thus be regarded as “cancelling out” across those policy options, in a sense.¹³⁰

This second approach—formally recognizing the existence of person-altering consequences but then explicitly ignoring them in the valuation calculations—is, however, also problematic. It is concededly common for conventional cost-

128. *See supra* note 109.

129. *Id.*

130. Jeffrey Gaba has argued that the very many different groups of persons whose conception and birth would be precluded by any policy chosen from the immense set of possible alternatives should perhaps have their interests ignored in an analysis on the basis that they “cancelled out” in the analysis. Gaba, *supra* note 93. While Gaba is only referring to cancelling out the interests of those groups of persons whose birth is precluded by the policy that is chosen, this cancelling out approach might be extended to also ignore, in an analysis, the interests of future persons who would be born as a result of the policy under consideration. Such a broader cancelling out approach with regard to both future persons who will be born as a result of a policy and potential future persons whose births are precluded by the policy has been suggested by Eric Posner:

For ordinary regulations such as environmental regulation, there will be little reason to think there is a morally significant difference between producing the first group of people [those born if the policy is not implemented] and producing the second group of people [those born if the policy is implemented]. Therefore I think the two would cancel out.

Email from Eric Posner, Kirkland and Ellis Professor of Law, University of Chicago Law School, to Gregory Crespi, Professor of Law, Dedman School of Law, Southern Methodist University (Feb. 5, 2008, 09:05 EST) (on file with the author).

benefit analyses to overlook impacts that the analyst cannot meaningfully reduce to quantitative terms. For example, the diffuse psychological impacts on persons due to their empathetic recognition of benefits conferred or costs imposed on other persons are difficult or impossible to measure with sufficient precision to meaningfully combine with more tangible policy impacts, and are commonly ignored. In many instances this practice of limiting the scope of the analysis to feasibly quantifiable impacts can be justified as a reasonable and necessary analytical simplification on the basis that those overlooked and practically immeasurable impacts are relatively small in magnitude relative to the more easily measured impacts, and/or they tend to offset one another in the aggregate, so the ultimate conclusions of the analysis are not significantly affected by their omission.

However, person-altering consequences are very different in these regards. First, they are huge in magnitude relative to the measurable impacts upon existing persons, and obviously cannot be overlooked without dramatically affecting the result of the analysis. Second, they are not internally offsetting; virtually all future persons would likely ascribe very large net benefits to a policy that is a necessary condition of their existence. The conventional arguments offered for selectively overlooking certain difficult to quantify policy impacts in a cost-benefit analysis, thus, do not support doing so when person-altering consequences are involved. Another shortcoming of this approach is, again, the downward bias problem noted above with regard to the offer prices/very high discount rate approach. Assigning zero value to the impacts upon future persons¹³¹ may, under some circumstances yield results that are even less accurate than the arbitrary conclusions under the conventional approach and could again lead to an even more pronounced bias towards radically present-oriented policies.

As a practical matter both approaches discussed above would lead to cost-benefit analysis policy recommendations that are based upon consideration of only the impacts upon existing persons¹³² and ignore any impacts upon future persons.¹³³ These approaches would tip the scales even more in favor of radically present-oriented policies that provide current benefits and whose adverse impacts primarily occur in distant future years—such as my ocean waste dumping policy hypothetical—and against policies that primarily result in current costs and distant future benefits—such as measures that would impose restrictions on fossil fuel use in an attempt to mitigate long-term global warming consequences—than does the current conventional approach.

Between these two alternative valuation approaches I marginally favor the second approach over the first approach as at least being the more candid of the two methods for circumventing the future effects valuation problem posed by

131. *See supra* note 110.

132. *See supra* note 109.

133. *See supra* note 110.

person-altering consequences. Either approach would move the analytical ball forward, at least marginally, by initially recognizing in principle the pervasiveness of person-altering consequences, and only then subsequently declining in one fashion or another to attempt to meaningfully quantify those consequences, rather than starting with the demonstrably false implicit assumption that such consequences do not occur and then calculating the policy impacts on future persons on that arbitrary basis. The more difficult question, of course, is whether either of these ploys are sufficient to rescue cost-benefit analysis from the difficulties posed by the problem of person-altering consequences.

Clearly the demonstrably false conventional implicit assumption that person-altering consequences do not occur at all leads to both arbitrary results and potentially radical undervaluation of policy impacts on future persons, and must be discarded. But is it sufficient to salvage the cost-benefit methodology to formally recognize that such person-altering consequences do indeed occur and are in fact of far greater significance than all other consequences combined for the affected future persons, but then to simply ignore them in one fashion or another on the basis that they cannot be meaningfully quantified? I think not. While it is perhaps a small positive step to nominally recognize the existence of person-altering consequences, any mode of analysis that then fails to attempt to meaningfully quantify the impacts of those consequences for comparative purposes cannot credibly be advanced as a comprehensive framework for guiding policy decisions. What you would get from such new valuation approaches, such as those I have just discussed, is what the noted cost-benefit scholar E.J. Mishan has colorfully labeled "horse and rabbit stew," an unappealing concoction which retains its equine flavor no matter how carefully the rabbit is chosen for its taste.¹³⁴ For each of these approaches, the "horse" would be an overwhelmingly strong orientation to favor radically present-oriented policies that generate benefits for existing persons, regardless of the magnitude of their later adverse impacts on future persons, and to disfavor any future-oriented policy that imposed net sacrifices on existing persons. To many persons this radical present-orientation would be just as unacceptable a feature of a policy evaluation framework as is the current practice of simply ignoring person-altering consequences through the use of demonstrably false baseline valuation assumptions.

Things do not look good for the continued viability of cost-benefit analysis. Moreover, there is a further difficulty that would have to be faced even if measuring techniques for quantifying person-altering consequences could somehow be much improved. Assume for a moment that it were possible to develop plausible willingness to pay-based estimates of the magnitude of the person-altering consequences of policies, of the very large benefits for distant future generations that would result from each of the policies under consideration, so that these policies could then be meaningfully compared and ranked in an

134. MISHAN, *supra* note 2, at 160-62.

inclusive manner that incorporated those consequences. A fundamental bias would still remain, this time a radically future-oriented bias. Even minor differences in the relative size of the very large estimates of the future net benefits for the various policy options, even if discounted at relatively (though not implausibly) high discount rates, would in all likelihood completely dominate any differences among the policies with regard to their net impacts on existing persons. The cost-benefit analysis recommendations would then be made solely on the basis of which policy option created the greatest benefits for future persons. The preferences of all existing persons taken together would have essentially no weight in the decision.

Making important policy decisions without regard to the preferences of any or even all existing persons is a bizarre and absurd idea. So even if the magnitude of the person-altering consequences of different policies could somehow be meaningfully measured and compared in cost-benefit analyses they would still provide unhelpful results and recommendations. This conclusion further highlights the weaknesses of the willingness to pay valuation criterion that underlies cost-benefit analysis with regard to assessing person-altering consequences. I am forced to conclude that there are no easy fixes; cost-benefit analysis is simply no longer a credible analytical approach once person-altering consequences are recognized. We have little choice but to develop new and broader normative criteria that also incorporate a broader range of secular but non-consequentialist ethical premises, or even overtly theistic premises, in a manner that will allow for meaningful recognition of the person-altering consequences of policies. By definition, however, any analytical approach that gave any weight at all to policy consequences apart from their effects on specific future persons as measured by willingness to pay principles could no longer be considered to be "cost-benefit analysis" as that phrase is conventionally understood.

3. Developing Alternative Criteria for Assessing the Impacts of Person-Altering Consequences

Let me offer a few preliminary thoughts as to how the significance of a policy's effects on future generations might be assessed in a manner that is not linked to the willingness to pay of specific individuals, but which still allows for quantitative aggregation with a willingness to pay-based assessment of the consequences for existing persons that has been derived through a conventional cost-benefit analysis. None of the alternative lines of inquiry that I will suggest appear to me to be particularly promising, but we nevertheless need to come up with some new assessment techniques, so more efforts along these or other lines are definitely called for.

I will not attempt to offer or comment upon any theistically-based valuation criteria. To the extent that one or more of the mainstream religious traditions provide sufficiently precise guidance as to how to ascertain and quantify the consequences of policies for future generations, apart from the estimated willing-

ness to pay to enjoy or avoid those consequences of specific future individuals, such approaches may provide acceptable policy assessment techniques for those persons who embrace those particular traditions. However, let me briefly note that theistic criteria may well not embrace the sharp dichotomy between policy impacts that will affect existing persons and those temporally more distant impacts that will affect only future persons. They may, therefore, call into question the use of willingness to pay-based valuations of the impacts of policies upon existing persons as part of an overall assessment methodology. I will here only address the possibility of developing new secular valuation criteria that differ from conventional cost-benefit valuations that are based upon the willingness to pay of specific individuals, and that could be used to value in quantitative terms policy impacts upon future persons and would thus allow for mathematical aggregation of those valuations with the conventional willingness to pay-based valuations of policy impacts upon existing persons.

As a threshold matter, it is clear that any attempt to value a policy's impacts upon future persons in a willingness to pay-based manner with reference to any hypothetical baseline scenario other than the actual alternative of those persons' non-existence would be subject to the same devastating criticisms made of the current practice of utilizing a demonstrably unattainable baseline reference scenario that assumes that those same persons would exist but would not experience the policy's impacts. Some other tact will have to be taken. There appear to be three possible lines of inquiry with regard to the development of a secular valuation criterion that is not grounded in their estimated willingness to pay of specific future persons.

One possibility would be to continue to focus upon policy impacts upon specific future persons, but value those impacts through some methodology other than those persons' willingness to pay to enjoy or avoid policy consequences. A second possibility would be to develop a different consequentialist valuation criterion that focused upon policy consequences other than the impacts upon specific future individuals. The third possibility would be to take an explicitly non-consequentialist approach, and somehow assign a numerical value to each policy option without regard to its consequences. Under each of these approaches, the valuation thereby derived would then have to be aggregated with the subject policy's willingness to pay-based valuations with regard to its impacts upon existing persons in order to obtain an overall policy assessment.

This first tact does not appear to be a very fruitful approach. Each policy option, because of its person-altering consequences, will lead to a different population of future individuals over a very extended period of time. The chosen criterion would have to be able to quantify the effects of each policy for each of the members of each of the relevant immense population groups. The same seemingly insurmountable problem of measurement uncertainty that would plague attempts to meaningfully compare the relative impacts of policies under the willingness to pay-based criterion, when person-altering consequences are

taken into account, would therefore exist here as well, unless that new valuation criterion, through its time discounting approach or otherwise, accorded only minimal significance to policy impacts upon future persons as compared to the impacts upon existing persons. But if it did so, it would again, as do the two willingness to pay-based alternatives that I have discussed in the prior subsection of this Article, consistently favor radically present-oriented policies. It would likely be widely regarded as unacceptable on that basis alone.

The second possible analytical move would be to attempt to value the impacts of policies on future generations by a different consequentialist criterion that does not consider impacts upon specific future persons as one of the relevant consequences. There is a modest literature that has begun to explore the use of such alternative consequentialist policy evaluation criteria, and I have summarized and discussed that work in an earlier article that focused on the ethical implications of the problem of person-altering consequences.¹³⁵ That literature, however, unfortunately has some rather severe shortcomings and limitations.

First, those writers have focused primarily upon establishing the ethical underpinnings that might justify recognizing such a consequentialist obligation to future generations that is not grounded in the impacts upon specific future individuals. They have not, however, yet attempted the far more difficult undertaking of proposing and defending a valuation algorithm by which the extent to which a policy implicates such an obligation could be quantified for aggregation with the willingness to pay-based impacts of the policy upon existing persons. On what possible basis could meaningful numbers be assigned to the discharge (or violation) of such impersonal obligations that do not run to specific individuals?

Second, and more importantly, that literature is unconvincing on its own

135. Crespi 2007, *supra* note 25, at 10881–84. Scholars who have made contributions to this inquiry include, among others, Anthony D’Amato, Edith Brown Weiss, Lothar Gundling, Michael Laudnor, William Grey, Rahul Kumar, and Doran Smolkin. *Id.* at 10883–84. One tact some of these writers have suggested is to define some sort of “rights” for future persons to not be burdened with overly adverse consequences, even if those consequences are logically necessary conditions of their existence, and then incorporate respect for those rights in some fashion in policy deliberations. *See, e.g.*, Bruhl, *supra* note 93; Smolkin, *supra* note 89. Another possible approach is to define various collective groups of persons as morally significant entities that have both rights and interests, such as specific future generations or even the “human race” as a whole, apart from the specific individuals that comprise those groups, and then assign values to impacts upon those entities. For arguments that have been made along these lines, *see, e.g.*, Kyser, *supra* note 7, at 37–38 (“[O]ne promising mechanism for doing so is to conceive of the ‘communities which future persons belong to [as] deserving of concern and respect in their own right.’”); Weiss, *supra* note 89, at 203–205 (referring to “group rights, as distinct from individual rights” that create a duty grounded in “planetary, or intergenerational rights [that] are not rights possessed by individuals.”); Gundling, *supra* note 89, at 207 (same); D’Amato, *supra* note 89, at 197–198 (“[i]t is somehow wrong to despoil the environment . . . even when we cannot calculate how such acts would make any present or future persons worse off.”). *See also* Michael Laudnor, *In Defense of Wrongful Life: Bringing Political Theory to the Defense of a Tort*, 62 *FORDHAM L. REV.* 1675, 1679–80 (1994); William Grey, *Possible Persons and the Problems of Posterity*, 5 *ENVTL. VALUES* 161, 168–172 (1996); Rahul Kumar, *Who Can Be Wronged?*, 31 *PHIL & PUB. AFF.* 99, 116–117 (2003).

terms even as to the threshold ethical criterion that it posits. Those writers generally argue in one fashion or another that such an obligation is grounded in an impersonal duty to the human race as a whole. These arguments in turn implicitly rest upon the view that the human race meaningfully exists in a moral sense apart from the individuals that comprise it.¹³⁶ However, the human race is comprised only of the specific individuals who come into existence over time, and to endorse a particular policy option as the preferred alternative on human race-enhancing grounds, so to speak, would simply serve to privilege one group of potential future persons who would be brought into existence under that policy over the vast multitude of other groups of persons who would be conceived and born under each of the virtually infinite number of policy alternatives. It is most unclear how this privileged group would be selected, and how the net benefits to this group would then be quantified under such a criterion for aggregation with the willingness to pay-based impacts of policies upon existing persons.

The third possible tact would be to abandon altogether the attempt to formulate a consequentialist valuation criterion, and simply ignore policy impacts on specific future persons for valuation purposes, and instead assess the extent to which each policy option conforms with the chosen non-consequentialist ethical criterion. Examples of such criteria would be the nature of the motives of the policymaker with regard to their degree of altruistic concern for the welfare of future generations, or the degree to which each policy option evidences a Kantian aversion to imposing uncompensated harms on future persons without their consent. Such a non-consequentialist valuation approach also appears to me to be unpromising, however, partly because of the obvious difficulties posed in reaching a consensus as to the appropriate criterion, and partly because I frankly cannot even imagine how such a criterion could then ever be quantified in application so as to render its valuations of a policy commensurate with the policy's willingness to pay-based impacts upon existing persons.

D. Conclusions Regarding the Problem of Endogeneity of Identity

Cost-benefit analysis is a widely used and highly influential analytical technique to evaluate policies on the basis of their consequences for affected persons, as measured by those persons' willingness to pay to enjoy or to avoid those consequences. The recent work of Derek Parfit and other philosophers, however, has made it clear that any social policy undertaken will trigger an exponentially cascading series of genetic alterations in subsequently-conceived persons that will eventually lead to the existence for all eternity of an entirely different group of future persons than those people that would have existed had the policy not been implemented. In other words, all policies have pervasive, eventually universal, and eternal person-altering consequences of overwhelming

136. Crespi 2007, *supra* note 25, at 10884 (citing extensively to the literature that takes this position).

significance relative to the other policy impacts. If cost-benefit analysis is to provide comprehensive and unbiased policy making guidance it will clearly need to take those consequences into account in some appropriate fashion.

Current cost-benefit analysts essentially ignore these person-altering consequences by implicitly assuming that policies will not have such consequences; that the hypothetical future persons whose valuations of policies are being estimated are the exact same persons who would exist absent the implementation of those policies. This demonstrably false assumption simply ignores the fact that these person-altering consequences are to the affected persons the most significant impact by far of the policies at issue, and any policy recommendations that are derived from incomplete and biased analyses that overlook these consequences are simply not relevant to the actual choices at hand. If one attempts to address this problem by incorporating person-altering consequences into cost-benefit analysis in the usual willingness to pay-based fashion that is the core premise of that approach, however, one reaches unhelpful results. All policies then generate truly massive future benefits of very uncertain magnitude, making it impossible to meaningfully discriminate among alternatives either with regard to their different future effects or with regard to the different trade-offs they present between current impacts and future effects.

There unfortunately does not appear to be an answer to this conundrum within the willingness to pay-based valuation framework of welfare economics upon which cost-benefit analysis is grounded. Continuing to ignore person-altering consequences by valuing the future effects of policies with regard to conceptually unattainable baseline scenarios that implicitly assume away those beneficial consequences to the affected persons will continue to lead to biased policy recommendations of a radically present-oriented character, and is therefore unacceptable. But there does not appear to be any way to meaningfully incorporate these person-altering consequences into the analysis and also avoid reaching the counterintuitive and unhelpful sweeping conclusion that any policy under consideration, whatever it may be, would result in massive benefits to future persons of uncertain magnitude that would greatly outweigh any adverse effects on existing persons. Attempts to avoid this result within the willingness to pay valuation framework through the use of discount rate manipulations, or through simply refusing to assign values to person-altering consequences, are evasive and unconvincing.

We appear to have little choice but to abandon cost-benefit analysis as it is now conducted. Moreover, any governmental decision-making approach whatsoever that limits itself to focusing upon policy consequences for specific persons, however those consequences are measured, is subject to essentially the same difficulties as is cost-benefit analysis with regards to assessing person-altering consequences.¹³⁷ We therefore need to try to develop broader normative criteria that are more inclusive in scope than merely considering policy consequences

137. I am indebted to Eric Posner for this insight. Posner, *supra* note 130.

for specific persons, and that will then allow analysts to meaningfully incorporate person-altering consequences into their assessments yet still provide policy recommendations that are both helpful and intuitively reasonable.

This will be a very difficult undertaking, to say the least, given the fundamental disagreements in this area. In particular, several very difficult questions are presented. Can a different consequentialist criterion be developed that focuses on consequences other than those that affect specific existing individuals? If not, then which if any non-consequentialist secular or theistic criteria might be appropriate to use as primary or at least supplementary normative standards for social policy making?¹³⁸ If some such alternative non-consequentialist criteria are to be used, then what quantitative weight should then be given to those criteria relative to the willingness to pay-based impacts upon specific individuals? These questions are indeed daunting, but the problem posed by person-altering consequences for policy analysis must somehow be addressed.

III. CONCLUSION

Cost-benefit analysis has long been subject to both external and internal criticisms. Notwithstanding these criticisms, it currently plays a crucial role in American public sector decision-making. Given this fact, it is important that such analyses be done as well as possible.

Prior critics of cost-benefit analysis have noted many of its shortcomings, but they have largely overlooked the fact that analysts implicitly assume that both the preferences of existing persons and the genetic identity of the members of future generations are exogenous, that is, unaffected by the policies that are being considered. This assumption is rarely if ever justified part of these analyses, and many analysts may even be unaware that their work rests on these unexamined premises. But as discussed above preferences are often to some significant extent endogenous with regard to the policies being pursued. If so, this leads to significant complications regarding the proper choice among the different preference structures to utilize for the valuation calculations. An unreflective utilization by an analyst of the pre-policy preference structures, under an erroneous exogenous preferences assumption that presumes that those preference structures persist indefinitely, may then lead to significant error.

In addition, the fundamental genetic identity of the members of future generations is invariably endogenous with regard to the policies that we pursue. A cost-benefit analysis that assumes that genetic identity is exogenous, and thus overlooks the pervasive and eventually universal and eternal person-altering consequences of the policy under consideration, will lead to results and recommendations that are simply irrelevant to the real choices at hand.

I have considered in this article both the endogenous preferences problem and

138. See *supra* note 135 and the accompanying text. I personally find these arguments to be most unconvincing in establishing meaningful additional criteria for evaluating policies that can be meshed in some plausible fashion with the assessments of policy effects on specific future persons.

the endogeneity of identity problem. With regard to the former problem, my main conclusion is that when preferences are to some extent endogenous with regard to the policy under consideration then cost-benefit analysis valuations should be derived solely from the post-policy implementation preferences, as a general matter. If there are any transitional preferences structures that precede the final, stable post-policy implementation preferences, then they should be used to value those policy consequences that occur while they are in existence. However, this theoretically appropriate valuation framework may need to be modified at times as discussed above to accommodate preference structure estimation difficulties, and to counter potential analyst bias, particularly with regard to cost-benefit analyses that are conducted prospectively rather than retrospectively.

The endogeneity of identity problem presents a much greater challenge. The current use of the demonstrably false implicit assumption that the genetic identity of the members of future generations is exogenous with regard to the policy under consideration is untenable. Since all policies have pervasive and eventually universal person-altering consequences, the results and recommendations derived from this exogeneity of identity assumption will massively understate the benefits to future generations as measured by those persons' hypothetical willingness to pay for policy consequences. But it does not appear that this problem can be meaningfully addressed within the secular, consequentialist framework of welfare economics and cost-benefit analysis that underlies the use of the willingness to pay-based valuation methodology.

It may therefore be necessary to replace or at least supplement the cost-benefit analysis methodology with a policy assessment approach for policy impacts affecting the members of future generations that rests upon non-consequential ethical premises, either of a secular or theistic character, and to develop a valuation algorithm that then can translate such a non-consequentialist assessment of the merits of a policy into a dollar figure that can be aggregated with a conventionally-derived willingness to pay-based assessment of that policy's consequences for existing persons. These are each daunting prospects, but the severe problem posed for cost-benefit analysis by the endogeneity of identity can no longer be overlooked and must somehow be addressed.

