Valuing Environmental Health Risk Reductions to Children

PROCEEDINGS OF

SESSION II: HOUSEHOLD DECISION MAKING

A WORKSHOP SPONSORED BY THE U.S. ENVIRONMENTAL PROTECTION AGENCY'S NATIONAL CENTER FOR ENVIRONMENTAL ECONOMICS (NCEE), NATIONAL CENTER FOR ENVIRONMENTAL RESEARCH (NCER), AND OFFICE OF CHILDREN'S HEALTH PROTECTION; AND THE UNIVERSITY OF CENTRAL FLORIDA

October 20-21, 2003
Washington Plaza Hotel
Washington, DC

Prepared by Alpha-Gamma Technologies, Inc.
4700 Falls of Neuse Road, Suite 350, Raleigh, NC 27609

ACKNOWLEDGEMENTS

This report has been prepared by Alpha-Gamma Technologies, Inc. with funding from the National Center for Environmental Economics (NCEE). Alpha-Gamma wishes to thank NCEE's Kelly Maguire and Project Officer Nicole Owens for their guidance and assistance throughout this project.

DISCLAIMER

These proceedings are being distributed in the interest of increasing public understanding and knowledge of the issues discussed at the workshop and have been prepared independently of the workshop. Although the proceedings have been funded in part by the United States Environmental Protection Agency under Contract No. 68-W-01-055 to Alpha-Gamma Technologies, Inc., the contents of this document may not necessarily reflect the views of the Agency and no official endorsement should be inferred.

TABLE OF CONTENTS

Session II: Household Decision Making

Benefit Cost Analysis and the Entanglements of Love

Theodore C. Bergstrom *

November 1, 2003

1 Introduction

Suppose that after careful thought, a parent reports that she is willing to pay up to $100 to save her child from one day of cold symptoms.[1] How can we use her answer and those of other parents to analyze the benefit of public projects that affect child health? Would it make sense to calculate the expected number of child days of cold symptoms that a project would prevent and to multiply this number by the average willingness to pay of a sample of parents?

But if a child has two parents, maybe we should value the child's health at the *sum* of the two parents' willingness to pay. Or, if the two parents' answers differ, perhaps we should use the maximum or perhaps the minimum of their two answers. And what of the child's own valuation on improved health? Why shouldn't we count that as well? To answer questions like these, we need to

*Aaron and Cherie Raznick Professor of Economics, University of California, Santa Barbara. This paper was prepared for a conference on Valuing Environmental Health Risk Reductions to Children, sponsored by the US EPA's National Center for Environmental Economics, the National Center for Environmental Research, and the University of Central Florida. I am grateful for helpful suggestions from Robin Jenkins of the EPA and Mark Agee of Penn State University.

[1] A series of carefully conducted studies have asked questions similar to this. Viscusi, Magat and Huber [25] interviewed individual parents who were asked their willingness to pay for hypothetical "safer" insecticides and toilet cleaners that would reduce health hazards for their children by specific amounts. Liu et al [16] found that a sample of Taiwanese mothers were willing to pay an average of about US $57 to avoid a cold for one of their children and about $37 to avoid a cold for themselves. (An earlier study by Alberini et al [1] asking Taiwanese adults about their willingness to pay to avoid colds for themselves found very similar values.) Dickie and Ulery [11] asked parents about their willingness to pay for the reduction of cold symptoms for a single child and for themselves. Dickie and Gerking [10] surveyed individual parents, asking their willingness to pay for a hypothetical sunscreen that reduced risk of skin cancer by a specified amount.

step back and examine the logic of benefit-cost analysis as applied to families in which some individuals care about the well-being of others.

2 What Can Benefit-Cost Analysis Do?

Without explicit instructions about how to compare one person's benefits with those of another, we can not expect benefit-cost analysis to tell us whether a public project should or should not be adopted. The best we can hope for from benefit-cost analysis is to learn whether a project is *potentially* Pareto improving. To see why this is so, it is useful to revisit a device found in Paul Samuelson's beautiful 1950 paper, "Evaluation of the Real National Income [23]. Samuelson introduces what he calls "the crucially important utility possibility function" to help us understand the logic of benefit cost analysis.

Consider a community with two people. A possible project would produce a public good valued by both persons. In order to produce it, community members will have to reduce their total consumption of private goods. If the project is not implemented, then alternative divisions of the total quantity of private goods will determine different utility distributions. These possible distributions generate a "utility possibility frontier" as described by the curve UP in Figure 1.

Figure 1: Utility Possibilities and Benefit Cost

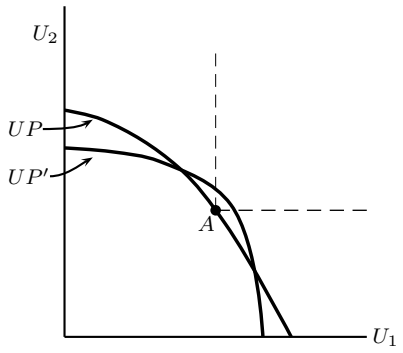

If the project is implemented there will be less private goods to be distributed between the two people, but each will enjoy a higher level of public goods. This will generate a new utility possibility frontier UP'. Because the curves UP' and UP are not "nested" there is not an obvious way to determine which is the better situation. Some utility distributions are attainable only if the project is not implemented and others are attainable only if the project is implemented.

Suppose that we know that we know that if the project is not implemented, then the utility allocation will be the point A in the graph. We see from the figure that the curve UP' contains points above and to the right of A. This implies that it is possible to produce the public good and to divide its costs in such a way that both individuals are better off than they were at A. When this is the case, we say that the project is potentially Pareto improving. Stating this formally:

Definition 1 *A public project is potentially Pareto improving if it is possible to implement the project and to distribute its costs in such a way that in the resulting allocation some community member(s) are made better off and none are made worse off.*

In a suitably simple economy, benefit-cost analysis can determine unambiguously whether a project is potentially Pareto improving. One such "suitably simple economy" is the following: A possible public project produces a public good that benefits several individuals and harms none.[2] The economy has a single pure private good.[3] Implementing the project requires the input of a known amount of private goods.

If individuals are selfish, the logic of benefit-cost analysis in this economy is very simple. An individual's "willingness-to-pay" for the project is defined to be the maximum amount of private good that she would be willing to sacrifice in order to enjoy the benefits of the public goods provided by the project. The benefit-cost test compares the sum of all individuals' willingnesses-to-pay for the project to its cost. It is quite easy to see that for this simple economy, the sum of individual willingnesses-to-pay exceeds total cost if and only if the project is potentially Pareto improving.[4]

The relation between benefit cost analysis and potential Pareto improvement becomes more complex when some of the beneficiaries are children who are loved and supported by parents, and even more complex if the parents care about each others' happiness. In this paper, we try to untangle the logic of these attachments a few strings at at time.

[2]It is not hard to extend these principles to an economy where some individuals benefit and others are harmed. But this simpler model suffices to illustrate the principles.

[3]The single-good assumption is appropriate in a multi-good economy if adopting and paying for the project does not result in changes in the relative prices of the private goods.

[4]If the sum of willingnesses to pay exceeds total cost, then it is possible to pay for the project while assessing no individual a share of the cost smaller than her willingness to pay. Implementing the project and paying for it with any such assessment constitutes a Pareto improvement. Conversely, if the project is potentially Pareto improving, there is a way to assess the costs of the project so that nobody is worse off after the project is implemented. Each individual's assessment would, by definition be smaller than his willingness-to-pay. Since these assessments add to the cost of the project it follows that the project would pass the benefit-cost test.

In order to extend the principles of benefit-cost analysis to a family, we need to specify and apply a theory of household decision-making for that family. Of course the world is filled with a great variety of family structures, which differ in membership and in the distribution of decision-making authority. As we will see, the prescriptions for benefit-cost analysis can vary significantly with the family type.

As we extend benefit-cost analysis to multi-person households, we face the question of how to reinterpret the criterion of potential Pareto improvement. This paper takes the position that the government is not able to intervene in the distribution of private goods within the households. Thus we consider a project to be potentially Pareto improving if and only if there is a way to assign the costs to families so that *given the household decision structure in each family*, no individual is made worse off and at least one is made better off.

3 A Single-Parent Household

Suppose that an economy is made up of households consisting of a single parent and a dependant child. Assume that each parent cares about her own consumption and her child's health and consumption and is selfish with respect to those outside her household. Assume also that children care only about their own consumption and health. Consider a government project that uses private goods as inputs and produces a public good which increases the health of children. How can we determine whether or not the project is potentially Pareto improving?

Let each parent's utility function take the form:

$$U\left(x, v(k,h)\right) \tag{1}$$

where x and k are the amounts of private goods consumed by the parent and the child respectively and h is a measure of the child's health. The child has no income and receives consumption goods only from its parent. The parent has an after-tax income, m which she allocates between her own consumption and that of her child.[5] Then, contingent on the child's health being h, the highest utility that a parent can achieve with after-tax income m is

$$U^*(m,h) = \max_{x+k\leq m} U\left(x, v(k,h)\right). \tag{2}$$

Consider a public policy that would increase her child's health from h to h'. We will define the parent's willingness-to-pay for the project to be W, where W is determined by the equation

$$U^*(m-W, h') = U^*(m,h) \tag{3}$$

[5] A more general model could allow the parent to allocate some private goods directly toward improvement of the child's health.

Definition 2 *If the sum of all parents' willingness-to-pay for a project exceeds the cost of the project, it is said to* pass the benefit-cost test for parents.

We can show that if a project fails the benefit-cost test for parents, then it is not potentially Pareto improving. To say this in another way:

Remark 1 *In an economy with single-parent households, if a project is potentially Pareto improving, then it must pass the benefit-cost test for parents.*

The proof is simple, but instructive. If the project is potentially Pareto improving, it must be possible to implement the project and to assign its costs it in some way so that no parent is worse off and at least one is better off than if the project is not implemented. Suppose that a project fails the benefit-cost test. Then in order for the project to be funded, some parent must be assessed a cost greater than her willingness to pay. Since the parents have complete control of their family incomes, this means that the parent who is forced to pay more than her willingness to pay for the project will be worse off if the project is implemented. Therefore the project cannot be potentially Pareto improving.

From Remark (1), we see that even though the Pareto criterion accounts for the well-being of children, it would be misguided to measure benefits by adding the value that children place on their own health to the values that their parents place on it. Adding children's valuations to those of their parents would lead to acceptance of projects that do not pass the benefit-cost test for parents and which are therefore *not* potentially Pareto improving.

The converse of Remark (1) is not in general true. That is, there may be projects that pass the benefit-cost criterion for parental preferences, but which do not allow a Pareto-improving reallocation in which the well-being of children as well as parents is (weakly) improved. This is not very surprising, especially since it has not been assumed that parents agree with their children about what is good for them. But, more remarkably, the converse fails to be true, even if parents and children share the same preferences over alternative combinations of child consumption and child health.

Parent's preferences are said to be *benevolent* if the parent's preferences over the child's consumptions coincide with child's.[6] Stated more formally, we say that

Definition 3 *A parent's preferences are benevolent if the parent's utility function is* $U_i\left(x_i, v_i(k_i, h_i)\right)$ *where* v_i *represents the child's preferences.*

Even if parents have benevolent preferences, it is possible that some projects that pass the benefit-cost test parental preferences will lead to outcomes that

[6] Though this usage is common in welfare economics, those who have been parents and those who have been children will recognize that parents who are "benevolent" by this definition would not always act in the best long-run interests of their children.

are worse for some children. This can happen because the child's health may be such a strong substitute for the child's income that with a healthier, the parent will reduce the amount of consumption goods transferred to the child by so much that the child is actually worse off after the transfer than before.

Figure 2: Parent's Consumption and Child's Utility

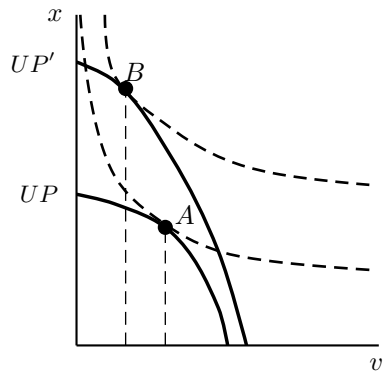

This effect is illustrated in Figure 3. The curves labeled UP and UP' shows combinations of parents' consumption and child's utility that are possible before and after a project is implemented. The point A on UP represents the preferred point of the parent if the project is not implemented and the point B on UP' represents the point preferred if the project is implemented. With the indifference curves that we have drawn, the parent will prefer B to A despite the fact that the child has a lower utility at B than at A. Notice that in Figure 3, the indifference curves drawn are consistent with parental benevolence, since the parents always prefer higher v to lower. One could imagine, for example, that for an impoverished family, an improvement in the child's health would induce the parents to require the child to earn its own food. Under these circumstances, despite their benevolence toward the child, the parents might favor an improvement in child health that is accompanied by such a large reduction in income transfers to the child that the child is worse off. Section 8.1 of the appendix, presents an algebraic example of preferences for which this is the case.

4 Lovebirds without Kids

Archie and Bess are a couple. They care about their own consumption of private goods and their own health. They also care about each other's happiness. Archie and Bess spend a good deal of time together and so each has become a good judge of the other's happiness, though this information arrives with a brief lag.

Let $x_A(t)$ and $x_B(t)$ be expenditures on consumption, and let $h_A(t)$ and $h_B(t)$ be measures of health for Archie and Bess respectively at time t.

Archie's happiness at time t is determined by the function

$$U_A(t) = v_A\left(x_A(t), h_A(t)\right) + aU_B(t-1) \tag{4}$$

and Bess's happiness is determined by the function

$$U_B(t) = v_B\left(x_B(t), h_B(t)\right) + bU_A(t-1). \tag{5}$$

Elsewhere [4] and [6], I show that this dynamical system is stable under plausible dynamics if and only if $ab < 1$. As these earlier papers explain, this stability condition limits the mutual intensity of their care for each other's happiness.

The government is considering a women's health project that will improve Bess's health by a known amount Δ. Archie and Bess are separately interviewed about their willingness to pay for this improvement. Bess realizes that if her health improves, this will have a direct effect on her happiness, which Archie will observe and enjoy, which in turn will make Bess herself happier, and so on *ad infinitum*. Can either Bess or Archie extract a reasonable answer to the interviewer's question from the blur of reflected happiness in this hall of mirrors?[7]

Equations (4) and (5) can be "inverted" to determine more conventional utility functions that depend only on the consumption and health of Archie and Bess. If consumptions and health levels are constant over time, equations (4) and (5) determine a time path of utilities for Archie and Bess that converges to equilibrium values. To find these equilibrium values, we set $U_A(t) = U_A(t-1)$, $U_B(t) = U_B(t-1)$ in Equations (4) and (5) and solve. Then $U_A^*(x_A, x_B, h_A, h_B)$ and $U_B^*(x_A, x_B, h_A, h_B)$ are the equilibrium utilities that result from the constant outcome (x_A, x_B, h_A, h_B). These utility functions are as follows:

$$U_A^*(x_A, x_B, h_A, h_B) = \frac{1}{1-ab}\left(v_A(x_A, h_A) + av_B(x_B, h_B)\right) \tag{6}$$

and

$$U_B^*(x_A, x_B, h_A, h_B) = \frac{1}{1-ab}\left(v_B(x_B, h_B) + bv_A(x_A, h_A)\right) \tag{7}$$

Let us call $v_A(\cdot)$ and $v_B(\cdot)$ the *private utility functions* of Archie and Bess respectively and let us call $U_A^*(\cdot)$ and $U_B^*(\cdot)$ their *social utility functions*. We assume that Archie and Bess are able to untangle their affections so as to find private and social utility functions as in Equations (6) and (7). There remain some tricky questions to answer. If know that their social utility functions

[7] Miles Kimball [15] suggested the hall-of-mirrors metaphor.

take this form, how can we use Archie's and Bess's reported valuations for improvements in Bess's health in a benefit-cost analysis? To answer this question, we need some assumptions about the way that Archie and Bess make family decisions.

4.1 Archie the Dictator

Gary Becker's famous "Rotten Kid Theorem" [2] operates on the assumption that family decisions are determined according to the preferences of a single benevolent household member.[8] Although this assumption is politically incorrect and, I believe, descriptively inaccurate for modern Western households, it remains worth studying for at least two reasons. First, the simplicity of this model makes it an instructive place to develop intuition and understanding that can be extended to other environments. Indeed since almost all economists interested in the economics of the family have cut their teeth on Gary Becker's patriarchal family [2], this model has become a benchmark against which alternative theories need to be tested. More importantly, not all of the households that interest us are modern and Western. Some of the most important applications of benefit cost analysis to public health programs are likely to concern traditional societies and highly patriarchal societies.

Assume that Archie has no direct control over his own health or that of Bess, but he is able to determine the amount of consumption expenditure that each receives, subject to a budget constraint. Then he chooses x_A and x_B to maximize his social utility function

$$U_A(x_A, x_B, h_A, h_B) = v_A(x_A, h_A) + av_B(x_B, h_B) \tag{8}$$

subject to the budget constraint $x_A + x_B = M$, where M is the family income. For given health levels and family income, the possible distributions of the private utilities, v_A and v_B are points lying on or below the curved line in Figure 4.1. Archie's indifference curves over these possible distributions will be straight lines with a slope of $-1/a$ and his preferred distribution of private utilities will be at the point marked A^*

Suppose that benefit-cost analysts interview Bess and ask her the following questions.

B.1 What is the largest amount of your personal consumption that you would be willing to give up in order to improve your health by Δ units?

[8] In Becker's treatment, although Archie is benevolent, Bess is selfish. I think he chose to do this not out of misogyny, but because he wanted to avoid the hall-of-mirrors complications of mutual affection.

Figure 3: Utility Possibilities

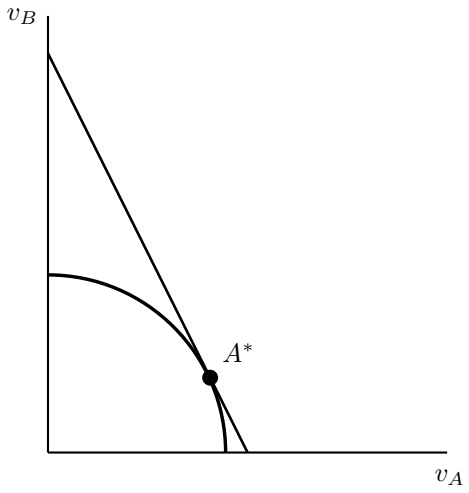

B.2 Given the way that Archie allocates income in your family, what is the largest amount of family income that you would want to give up in order to improve your health by Δ units?

Let us also suppose that they interview Archie and ask him:

A.1 What is the largest amount of family income that you would be willing to give up in order to improve Bess's health by Δ units?

When will the answers to these questions be different and when will they be the same? And how should we use the answers in benefit-cost analysis?

A Harmonious Case

Let us assume that Archie and Bess both have positive but diminishing marginal utility[9] for their own consumption of private goods and that Bess's private utility function is of the special form,

$$v_B(x_B, h_B) = f_B(x_B + kh_B) \tag{9}$$

[9]We have assumed that utility functions are additively separable between Archie's and Bess's private utilities. In the additive form, the v functions are unique up to positive affine transformations. The property of diminishing marginal utility for private consumption therefore has operational content in that it is preserved under all positive affine transformations.

where k is a constant and where the function $f_B(\cdot)$ displays positive but diminishing marginal utility. In this case Archie and Bess will be in perfect agreement about the value of an improvement in Bess's health. Bess's answer to question B.1 be will be the same as her answer to B.2 and also be the same as Archie's answer to question A.1.[10] In order to determine whether the health project is Pareto improving, a benefit-cost analyst could use this answer to represent the whole family's willingness to pay.

Remark 2 *If Archie and Bess both have diminishing marginal utility for private goods, if Bess's utility function is of the form* $v_B(x_B, h) = f_B(x_B + kh_B)$ *and Archie chooses the distribution of family income between private goods for himself and for Bess, then all three questions, B.1, B.2, and A.1 have the same answer, which is* $k\Delta$*. If the health project is implemented and Archie and Bess pay any amount less than* $k\Delta$*, both will be better off if the plan is implemented. If they must pay more than* $k\Delta$*,then if it is implemented, at least one of them will be worse off.*

It is easy to see that Bess's answer to question B.1 must be $k\Delta$. If she gives up $p < k\Delta$ of her own consumption and her health increases by Δ, the net effect is to increase her private utility and if she gives up exactly $p = k\Delta$ her private utility remains unchanged.

If the family spends p to improve Bess's health by Δ, Archie will allocate private consumption between himself and Bess by choosing x_A and x_B to maximize his utility, $v_A(x_A, h_A)+af_B(x_B+kh_B+k\Delta)$, subject to the constraint that $x_A+x_B = M-p$. If we define $y_B = x_B+h_B+k\Delta$, then we can restate this maximization problem as the equivalent problem: Maximize $v_A(x_A, h_A) + af_B(y_B)$ subject to $x_B + y_B = M + kh + (k\Delta - p)$. If the project is implemented and the family pays p, the net effect is seen to be the same as that of increasing Archie's budget by $k\Delta - p$. Thus Archie would be willing to pay any amount p up to $k\Delta$ for the project. Therefore Archie's answer to Question $A.1$ is $k\Delta$, which is the same as Bess's answer to $B.1$. Since we have assumed that both spouses have diminishing marginal utility of private consumption, it also true that an increase (decrease) in Archie's budget results in an increase (decrease) in both x_A and y_B and hence in Bess's private utility. Therefore Bess's answer to Question B.2 is also $k\Delta$.

Generous Archie?

With the utility function described in the previous section, Archie will agree to spend family funds on Bess's health if and only if she would be willing to pay for the increase entirely out of her own consumption. But if Bess's private

[10] This will be true for any positive a, that is, even if Archie is very selfish, so long as he cares a little about Bess's well-being.

utility function takes the additively separable form, $v_B(x_B, h_B) = f_B(x_B) + kh_B$ where Archie and Bess have diminishing marginal utility of private consumption, Archie will be willing to "share the cost" of an improvement in Bess's health. In this case, Archie's answer to Question A.1 is greater than Bess's answer to Question B.1 and equal to Bess's answer to B.2.

More generally consider the case where Archie considers Bess's health and her consumption to be complements, where complementarity is defined as follows.

Definition 4 *Person A regards the health and consumption of Person B as complements if for any distribution of consumption between A and B, an increase in Person B's health does not reduce A's marginal rate of substitution between B's consumption and A's consumption.*

A proof of the following result can be found in the Appendix.

Remark 3 *If Archie and Bess both have diminishing marginal utility for private goods and if Archie regards Bess's health and consumption as complements, then Archie's willingness to pay for an improvement in Bess's health exceeds the amount that she would be willing to pay for this improvement out of her own consumption.*

In the example discussed in Remark 2, Bess's health and consumption are not complements but substitutes. In this case, if Bess's health improves, Archie reduces the amount of consumption goods that she receives so as to leave her private utility unchanged. It is possible however, even where Archie is benevolent, that an improvement in Bess's health would induce him to reduce her allocation of consumption goods by so much that she is left worse off than he would choose her to be if she were less healthy. This is a case where he regards her health and her consumption as very strong substitutes. The example that we discussed in Section 3, where an improvement in child health led the benevolent parent to choose an outcome that is worse for the child can be reinterpreted to illustrate this case, with Archie cast in the role of parent and Bess in the role of child.

5 Non-dictatorial Outcomes

In modern Western societies, we do not expect either adult member of an ordinary household to have dictatorial power over the allocation of private goods. We consider two alternate theories of household allocation. One of these is a bargaining model in which the Nash cooperative bargaining solution to the household environment. The other is a model in which Archie and Bess reach unanimous agreement on outcomes based on a shared notion of fairness.

5.1 Fairness Norms and a Household Welfare Function

Suppose that in the course of their relationship, Archie and Bess build a mutually agreed notion of "household fairness" which governs their decisions. As Samuelson [24] suggests, this consensus might be expressed by means of a household "social welfare function which takes into account the deservingness or ethical worths of the consumption levels of each of its members." We could reasonably assume that this social welfare function would be an increasing function of each member's social utility as defined in Equations (6) and (7). Thus the household welfare function would be of the form

$$W\big(U_A^*(x_A, x_B, h_A, h_B), U_B^*(x_A, x_B, h_A, h_B)\big). \tag{10}$$

An appealing case can also be made that the household welfare function should be linear in its two arguments, so that we have

$$W\big(U_A^*, U_B^*\big) = \alpha U_A^* + (1-\alpha)U_B^* \tag{11}$$

for some α between 0 and 1.

A social welfare function with linear weights on individual social preferences would be implied if the couple formed their ethical preferences by "veil-of-ignorance" reasoning of the sort: "How would you want to divide household resources if there were an equal chance that I were you and you were me?" Although this kind of ethical reasoning can establish the existence of a social welfare function of the form of Equation (11), it is not sufficient to determine the parameter α that weights Archie's U_A^* relative to Bess's U_B^*. For example, our argument so far gives one no special reason to think that $\alpha = 1/2$ is more "fair" than any other α between 0 and 1. The scaling of the functions v_A and v_B is arbitrary and if these function are rescaled, then in order to preserve the same household preferences, the weight α would have to be altered.[11] As they cooperate with each other, Archie and Bess would need to develop a household consensus that determines the weight α that both regard as "fair". As Binmore [8] suggests, in the long run we might expect these weights to be close to the weights that would sustain a Nash cooperative bargaining solution for the household, since if there were a significant and persistent difference, the partner who could do much better with bargaining and threatening might defect from the social consensus expressed by the household social welfare function.

Recall from Equations (6) and (7) that the social utility functions are themselves linear combinations of the private utility functions, $v_A(x_A, h_A)$ and $v_B(x_B, h_B)$. In particular, we have

[11] Abram Bergson [3] introduced the notion of a social welfare function for an entire society. Harsanyi [14] proposed ethical foundations for such a function. Binmore [8] suggested that small, tightly-woven groups like families, may evolve a concept of group fairness that is well described by a household social welfare function.

$$W\big(U_A^*, U_B^*\big) = \frac{\alpha}{1-ab}\Big(v_A(\cdot) + av_B(\cdot)\Big) + \frac{1-\alpha}{1-ab}\Big(v_B(\cdot) + bv_A(\cdot)\Big) \tag{12}$$

Since preferences determine utility functions only up to an increasing transformation, and since by assumption $ab < 1$, we have an equivalent household welfare function if we multiply the expression (12) by $1 - ab$. Doing this and combining terms, we have

$$W = (\alpha + b(1-\alpha))\, v_A(x_A, h_A) + (1 - \alpha + a\alpha)v_B(x_B, h_B). \tag{13}$$

If Archie and Bess share a household social welfare function, then the interpretation of their answers to benefit-cost questions is really easy. If money is taken from household income to help pay for a public project that improves Bess's health, they both agree about how the cost will be shared in terms of their private consumption. Thus each of them would give the same answer to the question "How much would you be willing to pay for an increase of Δ in Bess's health?" The appropriate number for benefit-cost analysts to use is the number given by either of them. It would be a mistake to use the sum of their answers, since they both would agree that the sum as twice as much as they would want to pay.

5.2 Bargaining Equilibrium

The pioneering work on models of household bargaining was done by Marilyn Manser and Murray Brown [19] and Marjorie McElroy and Mary Horney [20], who proposed to model household decision making with the Nash cooperative bargaining model. In these papers, a marriage is modelled as a static bilateral monopoly. A married couple can either remain married or they can divorce and live singly. There is a convex utility possibility set S containing all utility distributions (U_1, U_2) that could possibly be achieved if they remain married. The utility of person i if he or she divorces and lives singly is given by V_i. It is assumed that there are potential gains to marriage, which means that the there are utility distributions (U_1, U_2) in S that strictly dominate the utility distribution (V_1, V_2). These papers propose that the outcome in a marriage will be the symmetric Nash bargaining solution where the "threat point" is dissolution of the marriage with both persons choosing to live singly. According to the Nash bargaining theory, the outcome in this household will be the utility distribution (U_1^*, U_2^*) that maximizes $(U_1 - V_1)(U_2 - V_2)$ on the utility possibility set S.[12] In this theory the outcome in a marriage is completely determined by

[12]This expression is sometimes known as the Nash product. John Nash [21] proposed a set of axioms for resolution of static two-person bargaining games such that the only outcomes that satisfy the axioms maximize the Nash product on the utility possibility set.

the utility possibility set and by the position of the threat point, (V_1, V_2). This theory has the interesting prediction that social changes that affect the utility of being single will affect the distribution of utility within the household and hence may change household spending patterns, even if they have no effect on the budget of the household, while changes in the apparent distribution of earned income within the household will have no effect on the distribution of utility in the household if they do not change the threat point from being single.

Shelly Lundberg and Robert Pollak [18], [17] propose an alternative Nash bargaining model. They suggest that for many marriages the relevant threat point for the Nash bargaining solution should be not divorce, but an "uncooperative marriage" in which spouses would revert a "division of labor based on socially recognized and sanctioned gender roles." The Lundberg-Pollak model makes predictions that differ significantly from the divorce-threat model. For example, in the Lundberg-Pollak model, if government child-allowances are paid to mothers rather than to fathers, the threat point shifts in favor of the mothers. Accordingly, the outcomes of cooperative bargaining within households are likely to be more favorable to women. By contrast, in the divorce-threat model, changing the nominal recipient of welfare payments when the couple is together will have no effect on the household outcomes if there is no change in the beneficiary in the event of a divorce.

To many married persons it seems unlikely that couples would resolve disagreements about ordinary household matters by negotiating under the threat of divorce. If one spouse rejects the other's proposed resolution to a household dispute, the expected outcome is not a divorce. More likely, there would be harsh words and burnt toast until another offer or counteroffer is made. It might even be that if the couple were to persist forever in inflicting small punishments upon each other, it would be *worse* both of them than a divorce. Nevertheless, the divorce threat may not be credible because divorce imposes large irrevocable costs on both parties, while a bargaining impasse need last only as long as the time between a rejected offer and acceptance of a counteroffer.

Nash's axiomatic approach to cooperative bargaining gives us no direct guidance about the appropriate threat points for bargaining in a marriage. But useful insight can be gained from Ariel Rubinstein's [22] model of noncooperative bargaining with alternating offers. Ken Binmore [7] extended the Rubinstein model to the case where each bargaining agent has access to an "outside option". Binmore's model is like the Rubinstein model, except that each person has the option of breaking off negotiations at any time and receiving a payoff of the value of an outside option.

The Rubinstein-Binmore model, as applied to marriage lends formal support to the Lundberg-Pollak notion of bargaining. The non-cooperative bargaining model predicts that household outcome will either be the Nash bargaining so-

lution in which the threat point is delayed agreement and burnt toast so long as this Nash solution is better for both than divorce. Thus the divorce threat will influence the outcome only if the bargaining solution reached under the threat of burnt toast is worse than divorce for one of the partners. In case the divorce threat is relevant, one partner enjoys all of the surplus and the other is maintained at a point of indifference between being divorced and being single.

In our example with Archie and Bess, suppose that health levels are determined by public policy while the allocation of private goods is determined by bargaining. Suppose that the utility that each would obtain in the absence of agreement depends on total household income and on the health of each. Where x is total household income, let us represent these "threat point utilities" by the functions $T_A(x, h_A)$ and $T_B(x, h_B)$ The Nash bargaining theory predicts that the outcome of household bargaining will be an allocation of private goods that maximizes the Nash product

$$\big(U_A^* - T_A\big)\big(U_B^* - T_B\big). \tag{14}$$

subject to the constraint that $x_A + x_B = x$

For fixed levels of health, the Nash product in Expression (14) works like a utility function a household utility function in the sense that private goods are allocated in such a way as to maximize this function. If a public health project is implemented and their household is taxed, the utility possibility frontier for Archie and Bess will be altered and so will their threat points. The household bargaining theory predicts a new allocation of utility. Neither Archie nor Bess cares directly about the value of the Nash product, but only about his or her own utility in the solution to the bargaining problem. Therefore they will in general not agree on their willingnesses to pay for a specific change in health.

Figure 4: Health and Bargaining

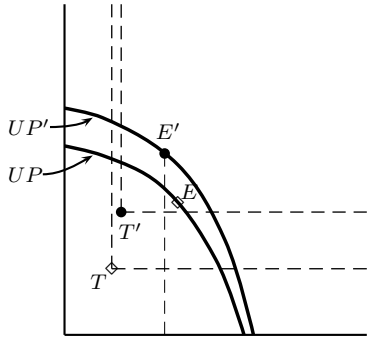

In a bargaining environment, not only is it possible that Archie and Bess have differing willingness to pay for a public project. It may even be impossible to get them to agree to fund a project that strictly increases the set of possible utility allocations.[13] This can happen because the project may shift bargaining power and/or the shape of the utility possibility frontier in such a way that the bargaining outcome is worse for one of them, even though the project shifts the utility possibility frontier strictly outward. This effect is is illustrated fairly in Figure 5.2. The initial utility possibility frontier is shown by the curve UP and the initial threat point is located at T. The point that maximizes the Nash product (14) is labelled E. Now suppose that a public health project is introduced and a fixed amount of tax is collected from the Archie-Bess household. One result of the project and the tax is to shift the utility possibility frontier outward from UP to UP'. A second result is to move the threat point from T to T'. Then the Nash bargaining theory predicts that the equilibrium outcome for the family moves from E to E'.

We see that a curious thing has happened. The health project and associated tax seems to be an unambiguous boon for Archie and Bess, since their utility possibility frontier moves everywhere outward. But if the project is implemented, the actual outcome in the household must be the corresponding bargaining solution E' and as we see in the diagram, E' assigns a smaller utility to Archie than his utility in the initial equilibrium E. Therefore Archie would oppose this change.

Thus if we continue to use the benefit-cost test to ask whether a public health project is potentially Pareto improving, we would have to use the *minimum* of Archie's and Bess's valuations to evaluate the benefits to this household.

6 Couples with Kids

After all this foreplay, it is time to bring a child into the Archie-Bess household and contemplate how to apply benefit-cost analysis to projects that improve children's health. A natural starting point for this theory is to think of the child's health as a "public good" that is valued by both Archie and Bess. Let us assume that Archie and Bess care about each others' happiness, as well as the health and well-being of their child and about their own consumptions of private goods.

Suppose that they have social utility functions of the form

$$U_A^*(x_A, x_B, x_C, h_C) = v_A\big(x_A, v_C(x_C, h_C)\big) + av_B\big(x_B, v_C(x_C, h_C)\big) \qquad (15)$$

[13] Lundberg and Pollak [17] apply similar reasoning to suggest that in the absence of binding contracts, married couples may not be able to make efficient household decisions about matters that affect relative bargaining power.

$$U_B^*(x_A, x_B, x_C, h_C) = v_B\big(x_B, v_C(x_C, h_C)\big) + bv_A\big(x_A, v_C(x_C h_C)\big) \qquad (16)$$

where x_C is the child's consumption, h_C is the child's health and where $v_C(x_C, h_C)$ measures the child's well-being.

6.1 Father Archie Calls the Shots

Let us begin by considering the Beckerian assumption that Archie is dictator of the allocation of private goods. Then given the level of child health h_C and family income x, Charlie chooses the consumption levels x_A, x_B, and x_C so as to maximize Expression 15 subject to the family budget constraint, $x_A+x_B+x_C = x$. Archie's willingness to pay for a small improvement Δ in the health of the child is approximately Δ times his marginal rate of substitution between child health and his own consumption. Assuming that he chooses to spend a positive amount of the household income on the consumption goods for Beth and for the child, Archie's marginal utility of consumption for the child is equal to his marginal utility of consumption for himself and also equal to his marginal utility of consumption for Bess. Therefore, Archie's marginal rate of substitution between child health and his own consumption is equal to his marginal rate of substitution between the child's health and the child's own consumption. But this marginal rate of substitution is seen to be simply the ratio of the two partial derivatives of the function $v_C(x_C, h_C)$ that measures the child's health. Thus Archie's willingness to pay for a small increase Δ in the health of the child is approximately

$$W_A = \Delta\left(\frac{\partial v_C(x_C, h_C)}{\partial x_C} \div \frac{\partial v_C(x_C, h_C)}{\partial h_C}\right). \qquad (17)$$

Thus Expression (17) is Archie's answer to the question "How much money would you be willing to pay for an increase of Δ in the health of your child?"

Suppose that the child has the same perception of its welfare as do its parents and its preferences are represented by the same function $v_C(x_C, h_C)$ that appears as an argument in the parents' utilities. If the child is questioned about the amount of consumption that it would be willing to forego for an increase of Δ in its health, its answer would be Δ times its marginal rate of substitution between health and consumption. We see from Equation (17) that this is precisely the same as Archie's marginal rate of substitution between the child's health and family income. It follows that the child's answer would be the same as Archie's.

We seem to have a kind of "Rotten Kid Theorem" here, but does it extend to Bess? Given that Archie dictates the household allocation of consumption goods, how will Bess answer the question "How much of the household income would you be willing to surrender to pay for an increase of Δ in the health of your child?" Bess realizes that if family funds are paid to support child health, the

resulting changes in private consumption by family members will be determined according to Archie's preferences. Given that Bess realizes this, the amount she would be willing to have the family spend to pay for an improvement of Δ in the child's health can be shown to be

$$W_B = W_A \left(\frac{\partial U_B^*}{\partial x_C} \div \frac{\partial U_B^*}{\partial x} \right). \tag{18}$$

where W_A is Archie's willingness and where $\frac{\partial U_B^*}{\partial x}$ is the marginal utility of household income to Beth, given her knowledge of how Archie will use it.

A simple calculation shows that

$$\frac{\partial U_B^*}{\partial x} = \frac{\partial U_B^*}{\partial x_A}\frac{dx_A}{dx} + \frac{\partial U_B^*}{\partial x_B}\frac{dx_A}{dx} + \frac{\partial U_B^*}{\partial x_C}\frac{dx_C}{dx}. \tag{19}$$

where dx_A/dx, dx_B/dx and dx_C/dx are the fractions of an incremental unit of income that Archie would allocate to the three household members.

From Equation (18) we see that Bess's answer to the benefit cost question will exceed Archie's answer if her marginal utility for the child's consumption is higher than her marginal utility for household income to be allocated according to Archie's priorities. This would be the case, if for example, Archie's distributional preferences lead him to allocate more private goods to himself and less to the child than Bess would like him to.

Which is the appropriate answer to use for benefit cost analysis? If we stick to the modest aim of identifying projects that lead to potential Pareto improvements and if Archie's willingness to pay is lower than Bess's, then we are forced to value the benefits at Archie's willingness to pay, which is the minimum of the two parents' answers. Using Bess's willingness to pay might also lead to a Pareto efficient outcome, but this outcome will not be Pareto superior to the initial situation, since Archie will be worse off.

6.2 Allocation with A Household Welfare Function

If household decisions are made by a shared household welfare function as discussed in Section 5.1, then as in our previous discussion, the situation is very simple. Archie and Bess would both agree in their rankings of possible alternatives for the family. Each would give the same answer to the question "How much household income would you be willing to spend for an improvement of Δ in your child's health. The answer given by either of them would represent the most that their household could pay in exchange for this gain without making them worse off.

6.3 Allocation with Divorced Couples

Suppose that after producing a child, Archie and Bess get a divorce. After the divorce, they have little affection for each other, but both care about the child's well-being. Bess has custody of the child and they set up separate household accounts, with each of them having a separate income. The only way that Archie can contribute consumption goods to the child is by giving money to Bess, who can allocate it as she chooses between herself and the child. Suppose that Archie makes no voluntary payments to Bess because only a fraction of the money that he gives her would be spent on the child's consumption.

In this situation, the appropriate measure for benefit cost purposes would be the sum of their two answers. If Archie is willing to pay an amount W_A and Beth an amount W_B for an improvement of Δ in the child's health, then both would benefit from implementing the project if amounts $P_A < W_A$ and $P_B < W_B$ are collected from the two parents. Thus their total willingness to pay is $P_A + P_B$. In this case, a benefit cost study that measured benefits of a child health project as the sum of the willingnesses to pay of a single parent for each child would underestimate the value of the project.

6.4 What About Grandparents and Aunts and Uncles?

A child's parents are not the only people who care strongly about its health and well being. Across almost all known societies, other close relatives besides parents are intensely concerned about the well-being of a child.

The theory of evolutionary biology gives us strong reason to expect that this must be true. The chances that our ancestors survived to reproduce successfully depended not only on the behavior of their own parents, but on the amount of help they got from their grandparents, uncles and aunts. If behavior toward one's near relatives is inherited, then the fact that we are products of successful evolution implies that humans can be expected to display significant willingness to care for their near relatives. In 1964, the great evolutionary biologist, William D. Hamilton [13] developed a theory that offers a systematic quantitative prediction of the extent to which individuals, on average, will care about their relatives. This prediction has come to be known as *Hamilton's Rule*.

Biologists define the "coefficient of relatedness" between two animals of the same species to be the probability that a rare "gene" found in one of these animals will also appear in the other. For sexual diploids if mating is between unrelated individuals, the coefficient of relatedness between two full siblings is 1/2, that between half-siblings is 1/4, that between an individual and a full sibling's child is 1/4, that between full cousins is 1/8, that between parent and offspring is 1/2, that between grandparent and grandchild is 1/4, and so on. Hamilton proposes that natural selection would produce individuals who try to

maximize *inclusive fitness* where inclusive fitness is defined to be a weighted sum of one's own reproductive success and that of one's siblings, half-siblings, and cousins of various types, where the weights are coefficients of relatedness.[14] (Of course some care has to be taken to avoid double-counting of say, one's children's fitness and that of one's grandchildren.)

Hamilton stated his rule as follow:

> The social behaviour of a species evolves in such a way that in each behaviour-evoking situation the individual will seem to value his neighbor's fitness against his own according to the coefficient of relatedness appropriate to the situation.

Hamilton's rule implies that when faced with the option of sacrificing c units of its own reproductive success in order to increase the success of a relative whose coefficient of relatedness is k, by b units, the decision maker should make the sacrifice if it passes the benefit-cost test $kb > c$.

Should we add the willingness to pay of a child's grandparents and other relatives to that of its parents? Or does this amount to double-counting as we found to be the case if we added the willingness to pay of its two parents in an intact family? To answer this question in a really satisfactory way, we would need a theory of decisions for extended families. In extended families where transfers occur or are expected to occur between the grandparents and the parents of a child, it might be argued that the family operates under a single budget and that a parent's answers to questions about her willingness to pay for improvements in the child's health incorporate her understanding of the psychic benefits this improvement will confer on the grandchildren and her beliefs about how the costs of this improvement would be shared under family transfers.

On the other hand, it seems likely that in most Western families, adult siblings have independent family budgets with no significant sidepayments between families. In this case it is appropriate for benefit cost analysis to add the willingness to pay of a child's other relatives to that of its parents. As we argued in the case of a divorced couple, if changes in a child's health and in the tax burdens of its relatives do not otherwise affect the pattern of transfers between relatives, then a public project that improves a child's health and costs its relatives less than the *sum* of their willingnesses to pay allows a potential Pareto improvement.

[14] More recent theoretical work [9] [12] [5] has shown that Hamilton's rule is strictly correct only where the benefits and costs from interaction between relatives take a linear form.

7 Conclusion

We have tried to untangle some of the strings of familial affection in order to give practical guidance for benefit cost analysis. Some of these strings untie relatively easily. For example, the members of a single-parent household will not enjoy a Pareto improvement if the household is assessed a cost greater than the parent's willingness to pay. Thus adding willingness to pay of children to that of their parents would lead a benefit cost analysis to accept projects that are not potentially Pareto improving.

We showed that with reasonable restrictions, the reflected happiness between loving couples can be resolved into a well-defined household welfare function over allocations. If the household acts to maximize such a function, then the willingness to pay stated by either parent is an appropriate measure of the benefits to the entire household. If there are relatives outside of the household who also have significant willingness to pay, then a proper accounting of benefits from child health should include the valuations of these relatives.

For households with bargained outcomes or for divorced households, interpretation of parents' responses to questions about their valuation of a child's health can, in principle be more complex. We have shown that married couples who allocate private goods by bargaining may differ in the amount of household funds that they would be willing to spend on the child's health. In these cases, strict adherence to the principle of identifying potentially Pareto improving projects suggests using the smaller of the answers proposed by the two parents. The recommendation is quite different for divorced couples so estranged that the noncustodial parent makes no voluntary contributions to the custodial parent's household. In this case, benefit cost analysis should add the willingness to pay of the noncustodial parent to that of the custodial parent.

8 Appendix

8.1 An Example

Each parent i has an income $m > 1$ and a utility function of the form $U_i = x_i + 2v_i(k_i, h_i)^{1/2}$, where the function $v_i(k_i, h_i)$ is a utility function representing child i's preferences. There are two possible values for h_i. These are $h_i = 0$ and $h_i = 1$. where initially $h_i = 0$, and the public project is adopted, $h_i = 1$, and where $v_i(k, 0) = k$ and $v_i(k, 1) = k/2 + 4/9$. The total cost of the project is $N/6$, where N is the number of parents. Each parent allocates her income x_i and k_i so as to maximize her utility. When $h_i = 0$, this occurs where $k = 1$ and $x_i = m_i - 1$. Thus if the project is not adopted, each parent i has utility $m_i + 2 - 1 = m_i + 1$ and each child has utility 1. If the project is adopted, then the parent the parent optimizes by choosing $k_i = 0$ and $x_i = m_i$. In

this case the utility of child i is $v_i = A = 4/9$ and the utility of the parent is $m_i + 2A^{1/2} = m_i + 4/3$. Each parent has a willingness to pay of $1/3$ for the project to be implemented and so the sum of willingnesses to pay is $N/3$. Since the cost of the project is $N/6$, it passes the parental benefit-cost test. But, as we have seen, if the project is implemented, each child's utility will be reduced from 1 to $4/9$.

Moreover, if the project is implemented, in order for children to be as well off as before the project, each child must receive at least $k = 10/9$. Any transfer scheme in which each child receives at least $10/9$ and where parents pay a total of $N/6$ for the project must be worse than the outcome without the project for at least some of the parents.

8.2 Proof of Remark 3

Let x_A and x_B be the consumptions that Archie would choose for himself and Bess if the the health project is not implemented. Let p_b be the maximum amount that Bess would be willing to sacrifice out of her own consumption in order to improve her health by Δ. Then it must be that Bess is indifferent between her initial situation and the allocation in which her health is $h_B + \Delta$ and her consumption is reduced to $x_B - p_B$. If the family budget is reduced by p_B and Bess's health is improved by Δ, then since Beth has diminishing marginal utility of private consumption and since Archie views Beth's consumption and her health are complements, it must be that Archie's preferred allocation of private goods out of the new budget is one in which Bess's consumption is reduced by less than p_B.

References

[1] A. Alberini, M. Cropper, T. Fu, A. Krupnick, D. Shaw, and W. Harrington. Valuing health effects of air pollution in developing countries: The case of Taiwan. *Journal of Environmental Economics and Management*, 34:107–126, 1997.

[2] Gary Becker. *A Treatise on the Family*. Harvard University Press, Cambridge, Ma, 1981.

[3] Abram Bergson. A reformulation of certain aspects of welfare economics. *Quarterly Journal of Economics*, 51:310–334, 1938.

[4] Theodore Bergstrom. Love and spaghetti, the opportunity cost of virtue. *Journal of Economic Perspectives*, 3(2):165–173, Spring 1989.

[5] Theodore Bergstrom. On the evolution of altruistic ethical rules for siblings. *American Economic Review*, 85(1):58–81, March 1995.

[6] Theodore Bergstrom. Systems of benevolent utility functions. *Journal of Public Economic Theory*, 1(2):71–100, 1999.

[7] Ken Binmore. Bargaining and coalitions. In Alvin Roth, editor, *Game Theoretic Models of Bargaining*, pages 259–304. Cambridge University Press, Cambridge, 1985.

[8] Ken Binmore. *Just Playing*. MIT Press, Cambridge, MA, 1998.

[9] L. Cavalli-Sforza, Luigi and Marcus W. Feldman. Darwinian selection and "altruism". *Theoretical Population Biology*, 14:268–280, 1978.

[10] Mark Dickie and Shelby Gerking. Valuing children's health: parental perspectives. Technical report, University of Central Florida, Dept of Economics, UCF, Orlando Fl 32826, 2003.

[11] Mark Dickie and Victoria Ulery. Parental altruism and the value of avoiding acute illness: Are kids worth more than parents. Technical report, University of Central Florida, Dept of Economics, UCF, Orlando Fl 32826, 2002.

[12] Alan Grafen. The hawk-dove game played between relatives. *Animal Behaviour*, 27(3):905–907, 1979.

[13] William D. Hamilton. The genetical evolution of social behavior, i and ii. *Journal of Theoretical Biology*, 7(1):1–52, July 1964.

[14] John Harsanyi. Cardinal welfare, individualistic ethics, and the interpersonal comparison of utility. *Journal of Political Economy*, 63:309–321, 1955.

[15] Miles Kimball. Making sense of two-sided altruism. *Journal of Monetary Economics*, 20:301–326, 1987.

[16] Jin-Tan Liu, James M. Hammitt, Jung-Der Wang, and Jin-Long Liu. Mother's willingness to pay for her own and her child's health: A contingent valuation study in Taiwan. *Health Economics*, 9:319–326, 2000.

[17] Shelley Lundberg and Robert Pollak. Bargaining in families. Technical Report 8642, NBER, 2001.

[18] Shelly Lundberg and Robert Pollak. Separate spheres bargaining and the marriage market. *Journal of Political Economy*, 101(6):998–1011, December 1993.

[19] Marilyn Manser and Murray Brown. Marriage and household decision theory–a bargaining analysis. *International Economic Review*, 21(1):21–34, February 1980.

[20] Marjorie McElroy and Mary Horney. Nash-bargained decisions: Towarda generalization of the theory of demand. *International Economic Review*, 22:333–349, 1981.

[21] John Nash. The bargaining problem. *Econometrica*, 18(1):155–162, January 1950.

[22] Ariel Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica*, 50(1):97–109, January 1982.

[23] Paul Samuelson. Evaluation of real national income. *Oxford Economic Papers*, 2(1):1–29, January 1950.

[24] Paul Samuelson. Social indifference curves. *Quarterly Journal of Economics*, 70(1):1–22, February 1956.

[25] W. Kip Viscusi, Wesley A. Magat, and Joel Huber. An investigation of the rationality of consumer valuations of multiple health risks. *RAND Journal of Economics*, 18(4):465–479, Winter 1987.

Preliminary research: Please do not cite

Weak Substitution, Environmental Vulnerability, and Choice
10/6/03

V. Kerry Smith, Mary F. Evans, H. Spencer Banzhaf, and Christine Poulos*

I. Introduction

Feenberg and Mills' [1980] concept of weak substitution provides the basis for expanding the role of revealed preference methods in measuring individual willingness to pay to reduce pollution-related health effects. To support this argument we outline a graphical analysis of how weak substitution can enhance our use of models based on averting and mitigating behaviors. Moreover, the proposed structure offers an economic framework for describing why different groups might be considered more or less vulnerable to the health effects of pollution. Our discussion begins by summarizing Smith and Banzhaf's [2003] proposed approach to describe how weak complementarity and the Willig condition contribute to the estimation of the economic value for reducing pollution's impact on environmental amenities.

Discussions of the preference restrictions used to recover information about consumers' values for non-market goods largely ignore weak substitution. For example, after reviewing the concept in some detail in his 1992 book, Freeman [2003] reflected the apparent professional consensus and excluded weak substitution from his revised edition. He maintains that an assumption of less than perfect substitution is not especially

* Paper prepared for presentation at EPA/UCF Workshop: Valuing Environmental Health Risk Reductions to Children, October 20-21, 2003. University Distinguished Professor, CEnREP, North Carolina State University and University Fellow Resources for the Future, Assistant Professor, University of Tennessee, Knoxville, Fellow, Resources for the Future, and Assistant Professor, University of Missouri, Columbia, respectively. Smith's research was partially supported by U.S. Environmental Protection Agency (Grant #R-82950801). Thanks are due to Dan Phaneuf for comments on aspects of this research and to Susan Hinton for help in preparing this paper.

informative without including a more detailed preference specification. Following Bartik [1988], Freeman argues that, at best, substitution offers upper and lower bounds for Hicksian welfare measures.[1] By contrast, our analysis suggests weak substitution may be comparable to weak complementarity in its ability to enhance the insights derived from revealed preference analysis of individual choice.

This paper provides a short primer on the logic underlying our arguments. We rely on graphical analysis and some simple analytical arguments but stop short of attempting to develop an empirical example. Section two outlines the first building block for understanding the link between weak substitution and choice. Our focus parallels Willig's [1978] use of weak complementarity. That is, we treat changes in environmental services as analogous to quality changes in private goods. This strategy redirects attention to selecting the private good whose price offers the most informative index of how a quality change affects consumer choice. Section three defines weak substitution and illustrates how a graphical analysis informs the structuring of economic models to recover preference information for the non-market resource with this restriction. By considering the price index capturing the effects of changes in environmental services, we isolate the private demands most likely to isolate an individual's non-market/market tradeoffs. Section four introduces a role for separability and with it a framework that demonstrates how an individual or household's economic circumstances, together with their physical (or health) conditions, contribute to a definition of their vulnerability to pollution. The last section describes how our logic relates to two past applications and

[1] Neill [1988] also discusses the prospects for substitution and complementarity relationships for groups of private commodities serving as bounds for willingness to pay measures.

how the framework might be used in new revealed preference (RP) or combined RP and stated preference research.

II. Analysis of Consumer Surplus with Indifference Curves

Following conventional discussions of Hicksian consumer surplus, we begin with the analysis of price changes. In analytical terms using indirect utility ($V(\cdot)$) or the expenditure function ($e(\cdot)$) with m designating individual (or household) income, p a vector of prices for market goods (X), q a non-market good, and u the utility level (e.g. $u^0 = V(m, p^0, q^0)$), the willingness to pay for a change in the price of X is defined implicitly in equation (1a) and explicitly in (1b).

$$V(p^1, m - WTP, q^0) = V(p^0, m, q^0) \tag{1a}$$

$$WTP = e(p^0, q^0, u^0) - e(p^1, q^0, u^0) \tag{1b}$$

Figure 1 provides the graphical interpretation that now finds its way into most undergraduate micro texts. Plotting the numeraire, z, on the vertical axis, a price reduction for good X from P_0 to P_1 corresponds to the pivot from AC to AD. WTP is the vertical distance AB.[2] This approach relies on changes in the prices of market goods and therefore has little bearing on non-market goods. As a result, many authors follow Lankford [1988] and adapt the early literature on rationing to deal with the WTP for non-market goods.[3]

In terms of expenditure functions, this strategy treats q as equivalent to a private good that is quasi-fixed. Letting r represent the price of q, we can define two expenditure

[2] It is labeled here with capital letters to distinguish it from the price vector.
[3] Freeman [2003] pp.74-81 has an excellent summary and graphical analysis.

functions– a conditional expenditure function that acknowledges the level of q can influence expenditures on X, labeled $e^*(\cdot)$, and an expenditure function which includes expenditures on q, labeled $e(\cdot)$. The following expression describes the relationship between $e^*(\cdot)$ and $e(\cdot)$:

$$e(p,q,r,u^0)=e^*(p,q,u^0)+r\cdot q \tag{2}$$

Notice that in equation (2) q remains quasi-fixed in defining $e(\cdot)$. If $\frac{\partial e}{\partial q}=0$ then we know the virtual price of q (i.e. $-\frac{\partial e^*}{\partial q}$) is equal to the exogenous price defined as r.

Figure 2 (adapted from Freeman's [2003] Figure 3.8) deals with this case where X is measured on the vertical axis and q on the horizontal axis. The WTP to realize a change in q from q^0 to q^1 is given in equation (3).

$$WTP=e^*(p,q^0,u^0)-e^*(p,q^1,u^0)-r(q^1-q^0) \tag{3}$$

Graphically, in Figure 2 the vertical coordinates of u^0 evaluated at q^0 and q^1 define WTP. The disparity in virtual expenditures (i.e. $r(q^1-q^0)$) explains the difference between WTP and $e^*(p,q^0,u^0)-e^*(p,q^1,u^0)$.

To this point our discussion is a "garden variety" explanation that can be found in Freeman or other treatments of the concepts relevant to Hicksian welfare measurement. The logic used to explain welfare measures often outlines how some factor, exogenous to individual choice, affects the expenditures required to maintain a given utility level. Frequently invoked assumptions include the existence of prices and the inability of individuals to adjust freely. That is, the individual is often assumed to pay a per unit

price for the good used to represent the environmental service. However, the choice process is described in a format that prevents her from actually choosing the amount desired at the available price.

As a rule, prices do not exist for environmental services. Thus, a comparison of Δe and Δe^* diverts attention from an issue that seems more relevant in models destined for empirical analysis. This issue concerns the selection of a private good (and hence a price) that provides the greatest insight into the importance of changes in q for an individual's observable choices. To develop this argument we need to consider a different way of describing the WTP for a change in q. This alternative entails selecting a single commodity and defining the price change that is *equivalent* to the change in q. We define this equivalence by specifying the changes in q and the selected good's price that require the same compensation (or payment) to restore the original utility level.

Thirty-five years ago, Fisher and Shell [1968] proposed a similar strategy in considering the issues associated with adjusting prices for quality changes. They demonstrate two results that seem underappreciated in the literature on non-market valuation.[4] The first of their results on price indexes relates to a quality change in a private good equivalent to augmenting the quantity of that good. This formulation of how quality influences preferences is the only one that allows the definition of a quality adjusted price (for the good experiencing the quality change) independent of the amounts of all other commodities being purchased. The second of their key theorems concerns another aspect of preference restrictions and focuses on conditions, other than simple repackaging (or equivalently quantity augmentation), that allow quality adjustment to the

price of one or more goods other than the one experiencing the quality change.[5] Their examples for this type of preference restriction implicitly recognize a role for weak complementarity[6]:

> "…suppose there is a quality change in refrigerators. If this change simply makes one new refrigerator deliver the services of some larger number of older ones, then the simplest price adjustment in the cost-of-living index is indeed an adjustment in the price of refrigerators. On the other hand, if that quality change also increases the enjoyment obtained from a quart of ice cream, then an adjustment in refrigerator price will not suffice; an adjustment in the price of ice cream is also called for. **Indeed, if the *only* effect of a refrigerator quality change is to augment the enjoyment attained from ice cream, then the simplest adjustment is one made *only* in the price of ice cream, even though the quality change takes place in refrigerators.** In this case, an adjustment in the price of refrigerators can be made to suffice; the magnitude of that adjustment, however, will depend on the quantities demanded of all goods. An adjustment in the price of ice cream will also suffice; the magnitude of that adjustment, however, will only depend on the quantity of ice cream and the quantity of refrigerators" (p.123, bold highlights added).

The conditions that assure the quality change can be reflected in a subset of private goods' prices parallel the conditions required to recover estimates of the role of non-market services (quality) from consumers' choices of private goods. This link is consistent with our proposal to consider the equivalence relation between quality changes and price changes as a guide to revealed preference modeling.[7]

To develop this logic we reconsider the measurement of consumer surplus for a price change, first in Marshallian terms. Consider one of the private goods, X, and a

---

[4] Banzhaf [2001] makes this point in discussing the adjustment of cost of living indexes for non-market environmental services.
[5] Hanemann [1984] addresses cases similar to this general idea in his evaluation of cross product repackaging for quality.
[6] This concept had not been defined at the time they developed their reasoning.
[7] As we noted, this argument is similar to Neill's [1988] use of substitution and complementarity restrictions for sets of private goods in relation to a non-marketed good to establish bounds for WTP.

change in the price of X from P_1 to P_0 (with $P_1 > P_0$). P_0TSRP_1 represents the change Marshallian consumer surplus (MCS) for an ordinary demand curve in panel a of Figure 3. Our argument is outlined for the linear case, but holds as an approximation for other cases. Equation (4) provides an algebraic description of P_0TSRP_1.

$$MCS = (P_1 - P_0)X_1 + \tfrac{1}{2}(P_1 - P_0)(X_0 - X_1) \tag{4}$$

Rearranging terms, we have equation (5):

$$MCS = \tfrac{1}{2}(P_1 - P_0)(X_1 + X_0) \tag{5}$$

Panel b of Figure 3 provides an alternative representation of this relationship. Assume the interior budget constraint corresponds to a price of X at P_1 and the exterior to P_0. The tangency at D corresponds to X_1 and at B to X_0. Marshallian consumer surplus for this linear case is exactly equal to the average of CD+BA. This is established directly once we recognize that, with numeraire good priced at unity, the following relations hold:

$$\begin{aligned} m_0 L &= P_1 X_1 \\ m_0 N &= P_0 X_1 \\ m_0 M &= P_0 X_0 \\ m_0 J &= P_1 X_0 \end{aligned} \tag{6}$$

As a result, CD=$m_0 L - m_0 N$ and AB=$m_0 J - m_0 M$. Substituting from equation (6) yields the relationship given in equation (7).

$$\begin{aligned} \tfrac{1}{2}(CD + AB) &= \tfrac{1}{2}\left[(P_1 - P_0)X_1 + (P_1 - P_0)X_0\right] \\ MCS &= \tfrac{1}{2}(P_1 - P_0)(X_1 + X_0) = \tfrac{1}{2}(CD + AB) \end{aligned} \tag{7}$$

To use this logic to illustrate the connection between changes in q and equivalent (in welfare terms) price changes, we assume weak complementarity. Let X and q be weak complements. The indifference curves for a given utility level with different levels

of q will intersect at R on the vertical axis (i.e. at zero consumption level for X) as in Figure 4, panel a. This feature follows from weak complementarity. That is, when there is no consumption of X the non-market good has no value to the consumer. We assume that $q_1 > q_0$. Quality improvements imply an inward fanning of indifference curves because the levels of the weak complement (X) and the numeraire (z) required to maintain the same level of utility decrease as q increases.

Now consider the definition of price change for X equivalent to the quality change in q. For the income level defined by the vertical intercept T in the figure, the budget constraints tangent to indifference curves labeled $V(q_0)$ and $V(q_1)$ at B and C, respectively, provide the equivalent price change. Moreover, the average of CD and AB in this case corresponds to the Hicksian willingness to pay for this price reduction (e.g. from the budget constraint tangent to $V(q_1)$ to the one tangent to $V(q_0)$) or *equivalently* the WTP for the increase in environmental services from q_0 to q_1. This conclusion follows because the utility level associated with these indifference curves is the same by construction (i.e. $V(q_0)=V(q_1)$). As in our earlier discussion of price changes, the argument holds exactly for the case of a linear Hicksian demand function and approximately with nonlinear demands.

Panel b of Figure 4 introduces a third indifference curve, $\bar{V}(q_1)$, corresponding to a quality level of q_1 but a different utility level. Comparisons of $V(q_1)$ and $\bar{V}(q_1)$ isolate the individual's choices holding income constant. We examine the case where prices correspond to the outer budget constraint. Therefore, $\frac{1}{2}(CD+EF)$ corresponds to

the Marshallian consumer surplus for the Hicksian equivalent price change and thus for the change in environmental services as well.

Smith and Banzhaf [2003] use this relationship (i.e. between $\frac{1}{2}(AB+CD)$ and $\frac{1}{2}(CD+EF)$) to illustrate how weak complementarity and the Willig [1978] condition can be combined to recover Hicksian consumer surplus measures for the amenity values provided by enhanced quality. Here we argue that Willig's focus on the relationship between the Marshallian consumer surplus for a quality change relative to the level of consumption of a private good (that is a weak complement to quality) and the price adjustment for that good with respect to quality is another way of suggesting that the conditions required for quality adjusted price indexes parallel those needed to define welfare measures for changes in quality.

III. Weak Substitution

Feenberg and Mills [1980] define weak substitution as another form of demand interdependency that arises for private goods and environmental services when the relative prices of those interdependent private goods are "high". In the case of weak complementarity, the choke price of the private good is crucial to isolating the marginal value of an increase in the related non-market environmental service (see Smith and Banzhaf [2003] for details). In the case of weak substitution, the level of consumption of a private good is defined, but the interaction between the private good and quality (or non-market services) implied by this definition is different. That is, we assume there exists a level of consumption for the private good (the weak substitute) above which improvements in the non-market environmental service have no value. For price levels

of the weak substitute above the level corresponding to this threshold, consumption of the weak substitute is below the threshold and quality improvements are valuable. Formally, if P_a defines the price that induces the threshold consumption level, X_a, as in equation (8), then the condition for weak substitution is given in equation (9).[8]

$$-\left.\frac{V_P}{V_m}\right|_{P=P_a} = X_a \tag{8}$$

$$\left.\frac{\partial V}{\partial q}\right|_{P=\hat{P}} = 0 \;\; \forall \;\; \hat{P} \leq P_a \tag{9}$$

Figure 5 illustrates the fanning indifference curves associated with weak substitution. As in our discussion of weak complementarity, the level of utility is constant across the curves with the inward fanning denoting quality improvements. Figure 6 now combines the format used in our earlier analysis of weak complementarity to illustrate how the quality change can be translated to an equivalent price change in the presence of weak substitution.

Consider first panel b in Figure 6. We begin the analysis by selecting a price for X, denoted P_a, that will lead to X_a when there is no consumption of z. The pivot in the budget constraint from the outer budget line tangent to $V(q_0)$ at B to the inner constraint defines a relative change in the price of z. The constraint tangent at B corresponds to the lower price, $P_z(q_0)$, for the numeraire good z as given in Figure 6a. The inner budget constraint, tangent at C, relates to the higher price, $P_z(q_1)$. In this case, the weighted (by P_a) average of AB+CD (measured along the X-axis – i.e. the weak substitute) corresponds to the Hicksian surplus as detailed in equations (10a) through (10c).

[8] An important point that we discuss in more detail below concerns the fact that P_a is in fact a function defined by the specified level of X_a and the prices of other goods and income.

$$AB = \frac{1}{P_a}\left(P_z(q_1) - P_z(q_0)\right)\cdot z_0 \tag{10a}$$

$$CD = \frac{1}{P_a}\left(P_z(q_1) - P_z(q_0)\right)\cdot z_1 \tag{10b}$$

$$\frac{1}{2}(AB+CD) = \frac{1}{2}\left(\frac{1}{P_a}\right)\cdot\left(P_z(q_1) - P_z(q_0)\right)\cdot(z_1+z_0) \tag{10c}$$

When P_a is normalized to unity, this relationship reduces to one that is comparable to our earlier description of the measurement of the Hicksian willingness to pay for a quality change. In this case however, we use the price of the numeraire and not the price for the private good serving as the weak substitute.

As in the case of our discussion of weak complementarity, introducing a third indifference curve, here at the level corresponding to a quality level of q_1 (labeled $\overline{V}(q_1)$ in Figure 6b), also yields the Marshallian measure of the quality change (or price change in z) as $\frac{1}{2}\cdot P_a\cdot(AB+FG)$ in the figure.

In some respects, our simple interpretation in Figure 6b may be misleading. The relative prices of X and z determine consumption levels of both goods. As we noted, Figure 6b assures that when z=0, consumption of X equals X_a, the threshold beyond which q no longer has value. The pivoting of the budget constraints at this point are intended to illustrate increases in the relative price of z in comparison to X, with the price for X remaining at P_a.[9] To some degree this representation is artificial because the diagram illustrates relative prices for a given income. Thus, P_a could be represented for a

[9] The actual slope of the budget constraint is $-\left(P_X/P_z\right)$. Increases in the price of z relative to P_a for X pivot the budget constraint inward.

different income level without pivoting at X_a. Figure 7 recasts the diagram starting the budget constraint at a position where $P < P_a$ when z=0. In this case, we also see that P_z can provide a price index for q. The combined results of Figures 6 and 7 suggests that the relevant condition is on relative prices. As Fisher and Shell suggest, we need to consider adjusting the price index of z and the level of consumption (or the price) for X. For the case of weak substitution (which was not explicitly discussed in Fisher and Shell's analysis), relative prices imply levels of consumption for X below X_a will be informative about the value an individual derives from changes in the amount of q. Our graphical illustration focuses on the price of the numeraire, z, to illustrate how this adjustment connects to the same Marshallian consumer surplus for a quality change discussed in the case of weak complementarity.

If we pivot panel b in Figure 6 and place X on the vertical axis with z on the horizontal, then the level of z corresponding to X_a (labeled as z_a in the figure) appears to correspond to what Smith and Banzhaf [2003] define as weak complementarity "at a point". This alternative interpretation of weak substitution provides a graphical illustration for Fisher and Shell's argument associated with improvements in the quality of refrigerators and ice cream. We can use the price of refrigerators to reflect quality changes (the price for the weak substitute in our example) but to do so requires that we incorporate consumption levels of all other goods. By contrast, when the effect is exclusively on ice cream we can use the price of ice cream and the consumption of refrigerators (our point X_a) to define a price index adjustment for quality. In this case we are implicitly assuming that for small levels of consumption of ice cream (i.e. levels at or below z_a) it is either not stored (and hence the quality improvement in the refrigerator is

not important) or weak complementarity is associated with larger amounts of consumption that might be associated with specialty desserts (e.g. ice cream cakes that must be refrigerated for some period time).

These examples have been quite simple. As a result, we relegated all other goods to the background without really explaining how this outcome was accomplished. In short, z was implicitly a composite of all other goods. Our graphical strategy sought to focus on the properties of each of the two goods' prices as candidates for the quality adjustment. In the process we implicitly assumed that one of them was an index for everything else.[10] One way to take account of other goods' prices in developing restrictions that use weak substitution imposes separability. This next step is discussed below.

IV. Separability and Environmental Vulnerability

To limit our attention to the relationships between z, X, and q and allow a more explicit treatment of the remaining private goods, we introduce a new argument in the preference function, r, which designates a composite of all other goods. P_r represents its price (or price index).[11] Separability of z, X, and q from r implies we can write the indirect utility function as in equation (11).

$$V = V\left(v\left(P, P_z, q, m_v\right),\ \mu\left(P_r, m - m_v\right)\right) \tag{11}$$

Separability also assures that consumption choices can be described as if they were undertaken as part of a process that decomposes the budget into components for each set

[10] This was not what Fisher and Shell intended in their discussion of the problem.
[11] r can be a vector of goods and P_r a vector of prices. We do not need to be more specific here because our objective is to illustrate how separability allows the role of P_r to be limited to its influence on income effects.

of separable goods. Thus, we can specify m_v as the expenditures on X and z, given levels of q. m_v will be a function of all goods' prices (P, P_z, and P_r) as well as q and m.[12] This strategy "focuses" the patterns of influence. That is, separability allows us to distinguish the roles of the prices of goods in the sub-function, P and P_z, from all other goods' prices, P_r. All prices continue to affect the conditional demand for X and z. However, those outside the sub-function enter through reallocation of income about separable sets of goods. With this specification, the influence of both P_r and m are observed through the income effect. Perhaps more importantly the effects of q on the demands for other goods are observed exclusively through the reallocation of income (e.g. m-m_v) and their income effects. There are no separate substitution effects attributable to q with other private goods.

Consider now the definition of conditional weak substitution given in equations (12a) and (12b). Equation (12a) defines the conditional demand for X, given the budget allocation between X and z and all other goods implied by m_v.

$$-\left.\frac{v_P}{v_{m_v}}\right|_{P=P_a} = X_a \tag{12a}$$

$$\left.\frac{\partial v}{\partial q}\right|_{P=\hat{P}} = 0 \;\forall\; \hat{P} \le P_a^* \tag{12b}$$

We can simply recast the analysis of weak substitution in the previous section in terms of the indifference curves corresponding to the separable sub-function ($v(\cdot)$) in the direct utility function. An average of the changes in expenditures on the mitigating good, z, for a given budget allocation, measures the conditional compensating variation (CCV) (see

[12] See Blackorby, Primont, and Russell [1978], Theorem 5.5, corollary 5.5.1 and their discussion on pp.277-279.

Hanemann and Morey [1992]). $(P_z(q_1)-P_z(q_0))z_0$ evaluates initial consumption of z at the old and the new prices with $P_z(q_0)\cdot z_0$ representing initial expenditures and $P_z(q_1)\cdot z_0$ an estimate of the expenditures required to purchase z_0 at the new prices. $(P_z(q_1)-P_z(q_0))z_1$ considers analogous expenditure increments at the new level of z. The average of these two terms (adjusted by the price of X) provides the CCV for the quality change. This formulation makes explicit the dependence of P_a^* on the allocation of income to the activities we hypothesize to be included in $v(\cdot)$.

Prices, demographic features, or any variables that may affect the allocation process can change how X_a influences the derived value of P_a^*. Thus, we can explicitly describe how demographic and physical traits contribute to characterizing individuals who are differentially vulnerable to pollution (or equivalently to declines in environmental services). The physical attribute of vulnerability corresponds to increased sensitivity to the level of q. Health status, age (e.g. elderly groups or young children), or other non-economic variables can be specified to lead to different levels of X_a for each subgroup.[13] Equally important, this structure describes how economic conditions contribute to the implied price, P_a^*, associated with realizing this threshold level of consumption of the weak substitute.

The separability restriction provides a convenient functional specification that allows this logic to be incorporated with conditional demand models. An observed pattern of vulnerability for particular groups under this definition arises through both

[13] Conventional definitions of weak complementarity assume that the threshold for quality having an effect is confined to zero consumption. In this context it is not possible to distinguish separate reasons for

economic and demographic (or other) factors. Because the two considerations influence observed responses, efforts to isolate sensitive groups and measure their responses must resolve a difficult identification problem. The analysis must take account of the economic consequences of the threshold defining vulnerability.

Interpreting the separability restriction as a means of introducing household models illustrates an alternative use of the framework. For example, Chiappori's [1988] collective household model imposes budget decomposition on the allocation of income in the household through his structuring of the household decision process. Combining this preference assumption with a further restriction that different members of the household consume different goods, we extend the use of weak separability to develop some specific insights into the use of averting or mitigating behavior models within a household setting. To illustrate this point, consider a two-person household where changes in q only affect individual one. Each household member consumes an exclusive private good, represented by X_1 and X_2 respectively. In addition, assume that, for individual one, the exclusive consumption good, X_1, serves as a weak substitute for q. Chiappori's efficient household model yields budget decomposition similar to what we described in equation (12). However, the interpretation of the decomposition in the household setting is different.[14] In Chiappori's model, the expenditure groups represent consumption by different people in the household, rather than simply different groups of goods consumed by the same person as in our equation (12). Equation (13) clarifies the distinction. Let $\ell^1(\cdot)$ designate the first individual's indirect utility function viewed as

---

environmental quality across individuals and use these physical or technical distinctions along with differences due to economic circumstances (since X=0 for all cases).

[14] See Smith and Van Houtven [2003] for a discussion of some of the welfare implications of the model.

the solution to Chiappori's efficient household. $\ell^1(\cdot)$ is a function of all goods' prices, income, and *q*. The left-hand side of equation (13) simply indicates that the indirect utility realized by individual one depends upon the prices for all the goods consumed by the household, *q*, and income. The right-hand side recognizes the exclusive goods consumed by each of the household members. As given here, individual one consumes the first good but not the second.

$$\ell^1(P_1, P_2, q, m) = V^1(P_1, q, s(P_1, P_2, q, m)) \tag{13}$$

$s(\cdot)$ represents the income available to individual one. The remainder, $m - s(\cdot)$, is the portion of household income available to individual two.

Equation (13) suggests two possibilities for defining the threshold level of commodity one implied by weak substitution, either in terms of $V^1(\cdot)$ or $\ell^1(\cdot)$. The central issue is whether the definition for X_{1a} implies that the income allocation is held constant. Equation (14a) uses Roy's identity for the Marshallian demand to define the price function for P_1 (e.g. labeled here as P_{1a}) that yields X_{1a} in terms of $V^1(\cdot)$.

$$-\left.\frac{V^1_{P_1}}{V^1_s}\right|_{P_1 = P_{1a}} = X_{1a} \tag{14a}$$

Equation (14b) defines weak substitution holding individual one's available income ($s(\cdot)$) constant.

$$\left.V^1_q\right|_{P_1 = \hat{P}_1} = 0 \quad \forall \quad \hat{P}_1 \le P_{1a} \tag{14b}$$

Return to equation (13) and consider a change in *q* in terms of the expression that defines utility for individual one without the explicit separability. This process results in the identity given in equation (15).

$$\ell^1_q = V^1_q + V^1_s \cdot s_q \tag{15}$$

Substituting from (14b) for $\hat{P}_1 \leq P_{1a}$ we have equation (16) for prices that satisfy this inequality.

$$\ell^1_q = 0 + V^1_s \cdot s_q \tag{16}$$

In other words, the model suggests that the tradeoffs we observe associated with individual one's choices in response to a change in *q* reveal any reallocation of income due to the change in *q*. This result holds provided we assume his (or her) income is held constant in defining P_{1a}.

In contrast, if we define weak substitution in terms of $\ell^1(\cdot)$, the model suggests that, for P_{1a} defined in this way, the incremental value of changes in *q* is counterbalanced by budget reallocation within the household. Equations (17a) or (17b) illustrate the result.

$$\ell^1_q = 0 = V^1_q + V^1_s s_q \tag{17a}$$

$$\frac{V^1_q}{V^1_s} = -s_q \tag{17b}$$

At this level of analysis we are unable to select one of these descriptions as correct. Our point is to direct attention to the types of household information that should be collected and, potentially, to illustrate how the absence of some types of information may confound the process of isolating the effects of *q*.

Perhaps the most direct insight from this set of modeling alternatives is that if the analyst is to have any chance of recovering individual preferences for *q*, then empirical analyses must include information on the sensitive member of the household, his or her

choices, and the responses of other household members to changes in some element of environmental quality or services. In the next section we discuss how the framework presented above relates to a few recent papers that sought to use household responses and specific sensitive populations to estimate individual willingness to pay for improvements in q.

V. Implications for Empirical Models

Weak complementarity has been *the* dominant restriction used in revealed preference approaches to measure the economic values for changes in non-market environmental services. More generally, complementarities between goods and services are important to the processes through which new goods (or services) generate improvements in individual well being. Bresnahan and Gordon [1997] developed this important insight in the introduction to their volume on incorporating new goods in cost of living measures. They used the example of artificial light to introduce their argument, noting that over the past century:

> "A series of new goods, such as whale oil for lamps, gaslight, and then the electric light bulb, rapidly lowered the costs of using artificial light, a commodity which is complementary to a wide variety of household and workplace activities. Thus, as artificial light grew cheaper, activities which had been economic only for short parts of the day spread to evening, activities confined to summer became year-round, and jobs became easier to perform" (p.3)

These types of complementarities have cascading effects – with cheap artificial light transforming the allocation of individual time throughout the day and creating a new mix of demands for goods and services – both new and old.

An important dimension of these effects is a set of discontinuities which was not emphasized by Bresnahan and Gordon and is central to Smith and Banzhaf's [2003] generalization ot weak complementarities. Relationships between goods and services may change substantially at different levels of consumption. Recognizing and using this information should be regarded as a preference restriction. Weak complementarity at a point and weak substitution are restrictions that provide a conceptual basis for developing these restrictions in ways that can be used in understanding individual choice.[15] That is, weak substitution's definition requires that we identify a point of discontinuity in the role for quality (or environmental services). As a result, it is adding more information than alternative specifications of substitution patterns (as might be associated with classes of preference functions, such as a constant elasticity specification).

To our knowledge, none of the empirical analyses of the costs of mitigating or averting behavior has imposed weak substitution. Moreover, while several papers have offered recommendations for future empirical analyses that follow the implications drawn from our combined models with separability and weak substitution, these studies do not suggest how the additional data would be used in estimation.[16] By specifying a formal structure, we believe the empirical implementation may be clearer. To illustrate this point we selected two applications that consider environmental impacts on family members – Agee and Crocker [1994] and Mansfield et al. [2002].

(a) Agee and Crocker [1994, 1996]

[15] We owe this observation to Dan Phaneuf who commented on the Smith-Banzhaf argument for weak complementarity by noting it created discontinuities and with them more information for revealed preference models.

[16] Shogren [2001] listed as his top two recommendations in a paper on valuing effects of environmental hazards on children's health: (a) "pay more attention to how decisions of intrahousehold resource allocation and distribution are made by caregivers; and (b) begin efforts to construct a systematic framework to help

Both Agee and Crocker papers investigate parental willingness to pay for risk information and therapy for children with differing confirmed levels of body lead. The focus of their analysis is on the decision to undertake chelation treatment given information about health risks, parental background (including education and labor market status), body lead levels, and household income. The authors are careful to include time costs in their estimates for the full costs of treatment (e.g. blood chelation). They consider average estimates for parental time commitment but do not have specific information on time allocations to the chelation activity or other activities for each family. The income, employment status, wage rates, family composition, and education levels are specific to each group.

Agee and Crocker find the chelation decision was related to parental education, family status (e.g. father present), lead levels, and measures of the full price. The results suggest more educated mothers and traditional families are less likely to choose the therapy. The authors suggest this outcome could imply a substitution of caregiving time for the therapy. Our framework would suggest that this behavioral response is exactly where one should expect to recover information about the economic value of reducing lead in the environment based on its impact on children. However, the information available to these authors provides no specific details on time allocation choices of parents, children's health status (that might serve as the threshold for a weak substitution effect), or reallocation of resources within the household. The authors correctly conjecture these issues are likely to be important but are unable to estimate a structural model without these details.

---

organize how we think about the interaction of environmental threats and the behavioral choices of society, caregiver and children…" (highlights of author's summary, pp.7-8).

(b) Mansfield et al. [2002]

Mansfield et al. [2002] propose an ambitious diary study of families with asthmatic children to investigate parental responses to high ozone warnings. The authors hypothesize that changing the mix of activities undertaken by children offers a behavioral response to changes in air pollution. They use a national sample including asthmatics and non-asthmatics during times when high ozone might be expected so it seems reasonable to anticipate observed responses. The authors' preliminary results focus on total hours spent outdoors and the fraction of the day spent outside. Specifically, for each activity in which the child participated parents were asked to whether the activity was undertaken "totally indoors", "mostly indoors", "half indoors and half outdoors", "mostly outdoors", or "totally outdoors". Based on the time diaries, the authors calculated the total number of hours spent in each of these categories. The results do not adhere to one's a priori expectations. On average, parents of asthmatic children indicated that their children spent more time mostly outdoors more often than parents of non-asthmatics. Of parents who reported their child's time was spent "totally outdoors", the average hours were higher among parents of non-asthmatic children. However, as in the "mostly outdoors" category, the differences across the two groups were not large and did not appear to be significantly different. Multivariate analysis confirms the simple analysis – ozone warnings did not have a differential effect for the two groups.

The weak substitution framework suggests a potential explanation for this finding: the need to separate the sample with children distinguished by a health threshold. This partition might reflect parental judgments about when high levels of ozone are likely to be problematic in ways that are comparable to our weak substitution relationship. Our

model also suggests that the household's capacity to reallocate resources among its members should affect the recoverability of the air quality effect. Our comparison of weak substitution defined as $\ell_q^1 = 0$ versus $V_q^1 = 0$ illustrated the importance of resource allocation. While our model focused on reallocation in terms of income, a reformulation to reflect time reallocation is a straightforward extension. The lesson for this analysis is, as Shogren suggested, greater attention to caregivers' time and monetary reallocations.

Overall, these two examples suggest that behaviors consistent with the weak substitution framework seem quite likely. What appears to be missing are data that include not only responses concerning the weak substitute, but other goods or services (or in the case of the household, the person) that a weak substitution model suggests is the best source for choice information.

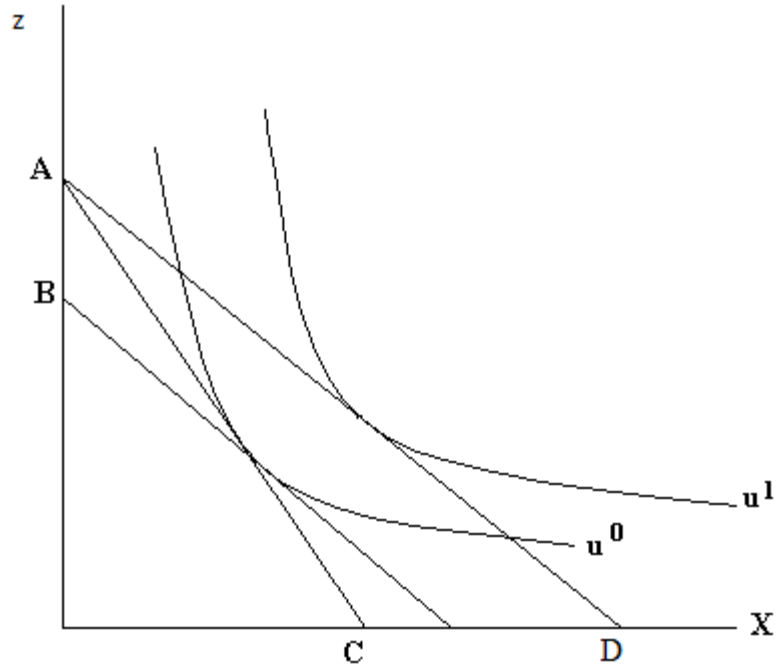

Figure 1: Conventional Description of Willingness to Pay – Price Reduction

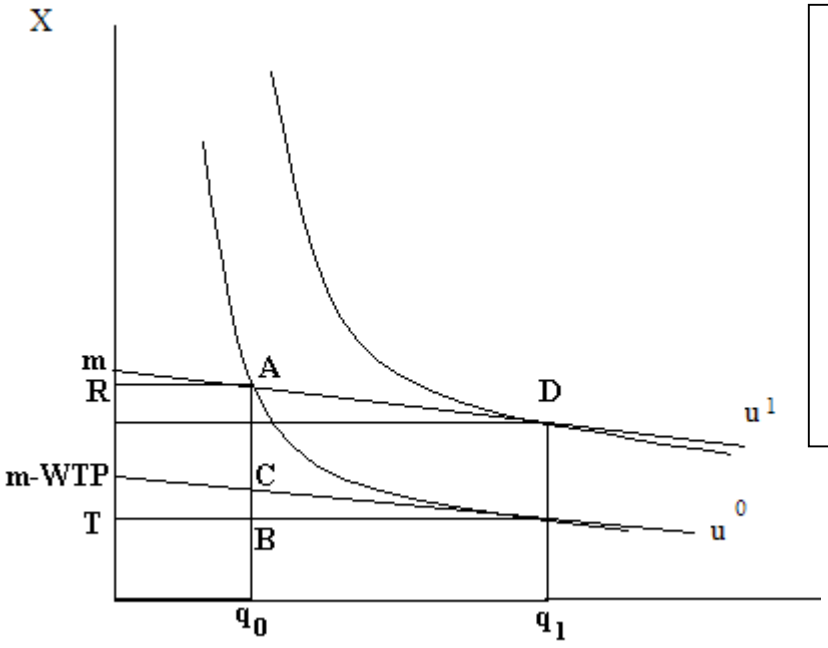

Legend for Figure 2

m = income

WTP = willingness to pay for a change in the rationed good from q_0 to q_1

$\Delta q = q_1 - q_0$

r = price of rationed good

A to $B = \Delta e^*$

B to $C = r \cdot \Delta q$

m to $R = r \cdot q_0$

Figure 2: Freeman's Description of Lankford Model – Change in Rationed Goods

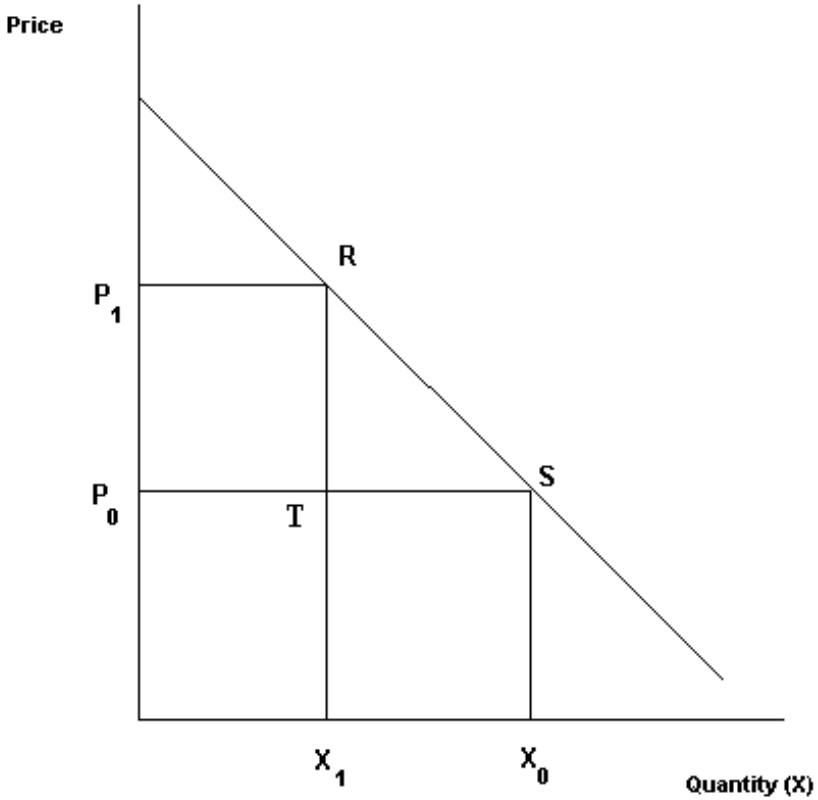

Panel a

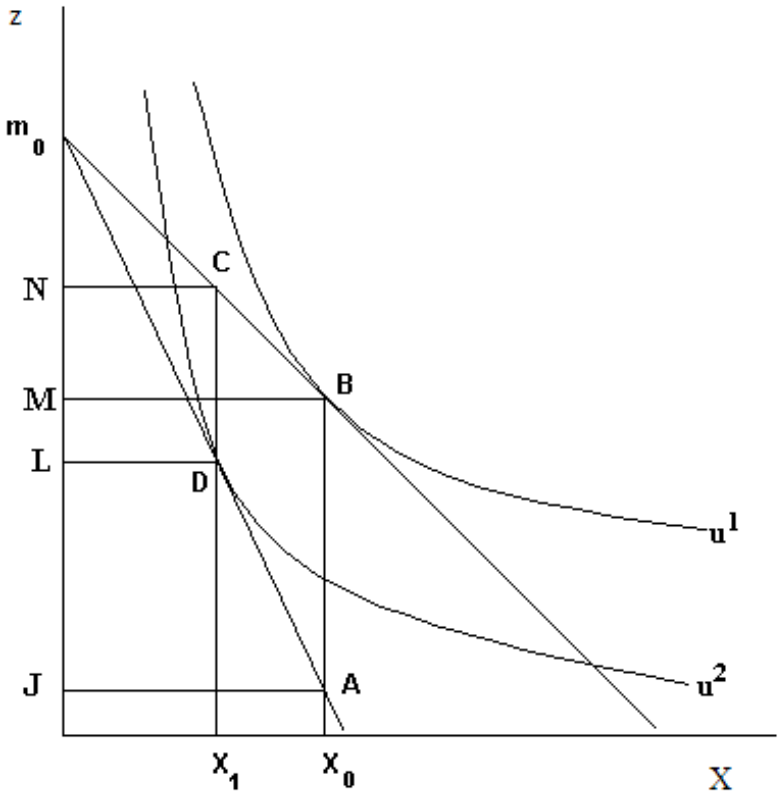

Panel b

Figure 3: Marshallian Consumer Surplus for a Price Change

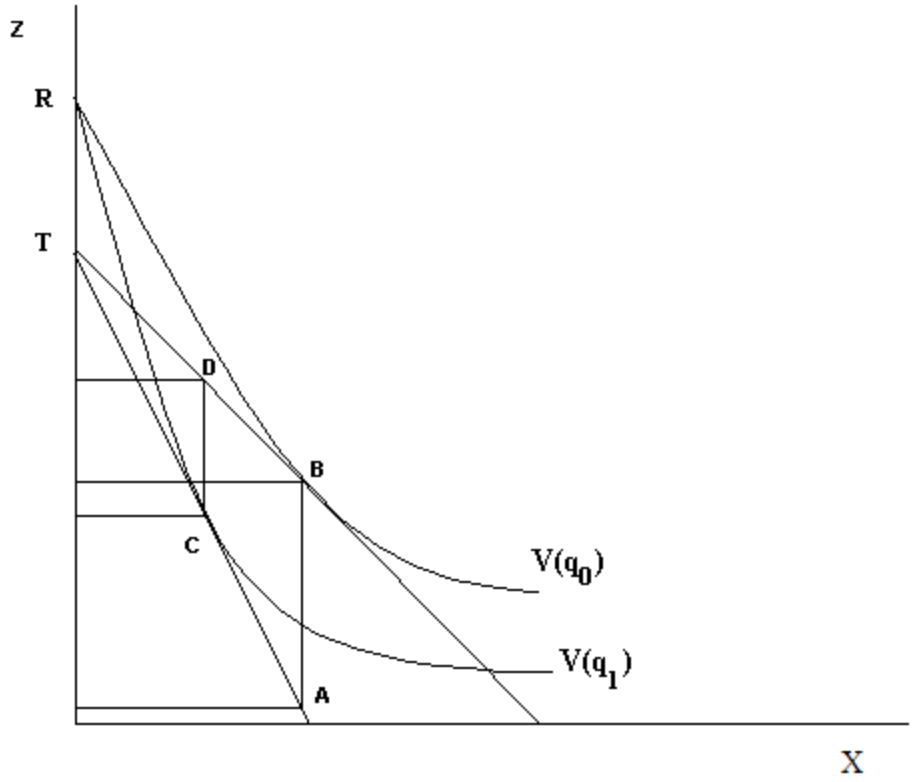

Panel a

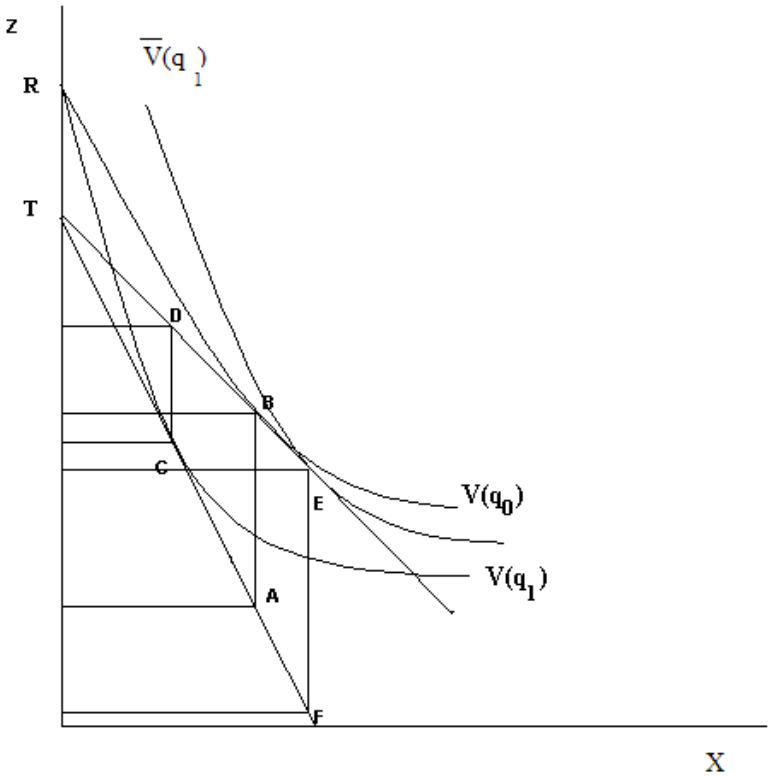

Panel b

Figure 4: Illustrating Quality Change

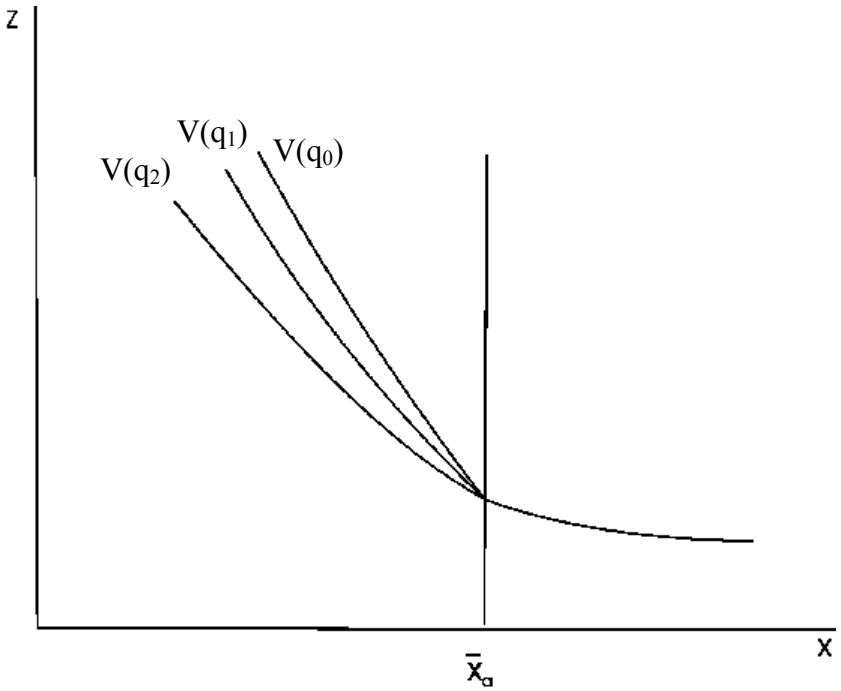

Figure 5: Weak Substitution

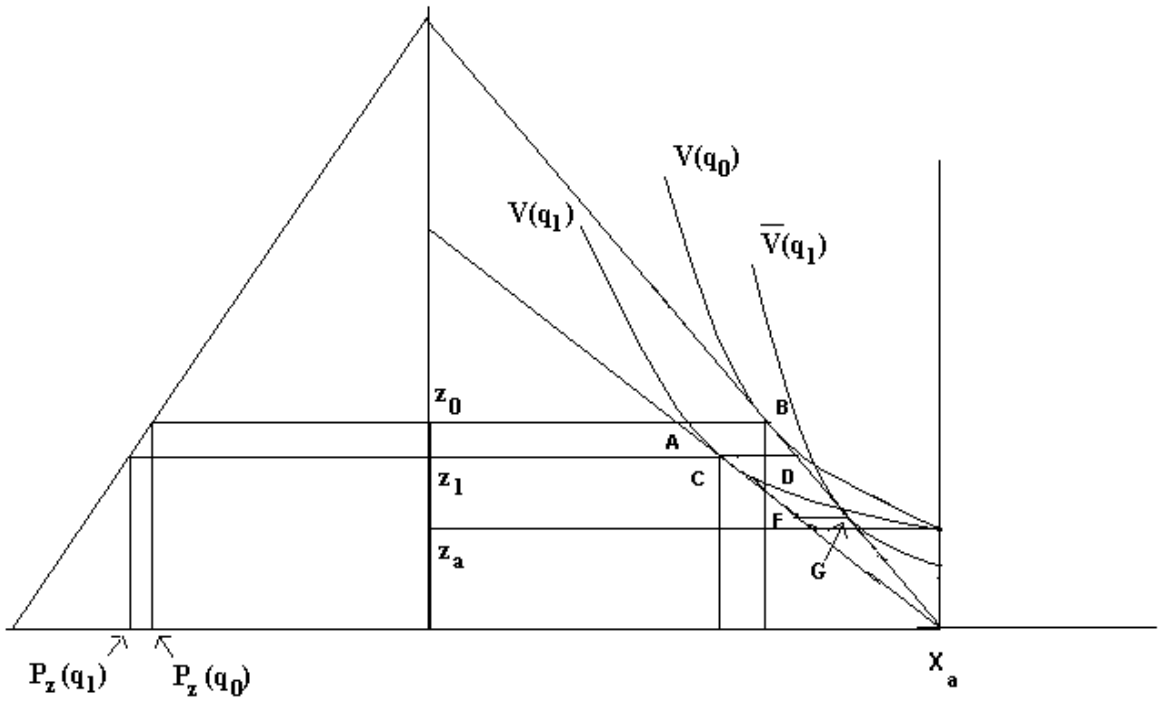

Panel a Panel b

Figure 6: Weak Substitution and Hicksian Demand

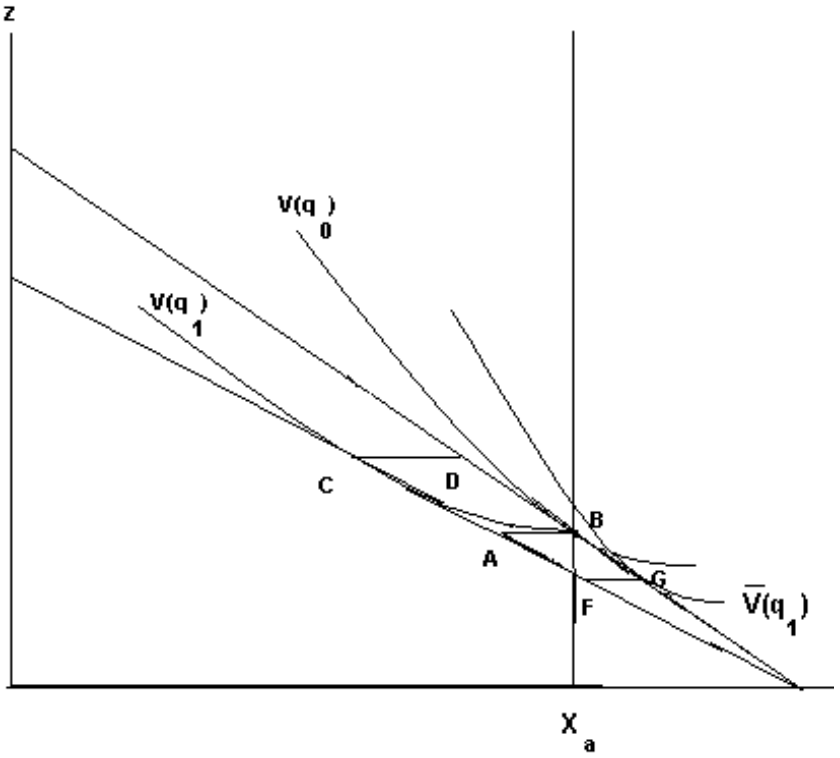

Figure 7: Weak Substitution with Alternative Relative Prices

References

Agee, Mark D. and Thomas D. Crocker, 1996, "Parental Altruism and Child Lead Exposure: Inferences from the Demand for Chelation Therapy," *The Journal of Human Resources* 31 (3): 677-691.

Agee, Mark D. and Thomas D. Crocker, 1994, "Parental and Social Valuations of Child Health Information," *Journal of Public Economics* 55 (September): 89-106.

Banzhaf, Spencer H., 2001, "Quality-Adjusted Prices: Theory, History, and Application to Air Quality," Ph.D. Dissertation, Department of Economics, Duke University.

Bartik, Timothy J, 1988, "Evaluating the Benefits of Non-marginal Reductions in Pollution Using Information on Defensive Expenditures," *Journal of Environmental Economics and Management* 15: 111-127.

Blackorby, Charles, Daniel Primont, and R. Robert Russell, 1978, *Duality, Separability and Functional Structure: Theory and Economic Applications* (New York: New Holland).

Blundell, Richard and Jean-Marc Robin, 2000, "Latent Separability: Grouping Goods without Weak Separability," *Econometrica* 68: 53-84.

Bockstael, Nancy E. and Kenneth E. McConnell, 1993, "Public Goods as Characteristics of Non-Market Commodities." *Economic Journal* 103 (September): 1244-1257.

Bradford, David F. and Gregory G. Hildebrandt, 1977, "Observable Preferences for Public Goods," *Journal of Public Economics* 8 (October): 111-132.

Bresnahan, Timothy J. and Robert J. Gordon, 1997, "Introduction," *The Economics of New Goods*, edited by T.J. Bresnahan and R.J. Gordon (Chicago: University of Chicago Press for NBER).

Chiappori, Pierre-André, 1988, "Rational Household Labor Supply," *Econometrica* 56 (January): 63-89.

Feenberg, Daniel and Edwin S. Mills, 1980, *Measuring the Benefits of Water Pollution Abatement* (New York: Academic Press).

Fisher, F.M., and K. Shell, 1968, "Taste and Quality Change in the Pure Theory of the True Cost-of-Living Index," in J.N. Wolfe, editor, *Value, Capital and Growth: Papers in Honor of Sir John Hicks* (Chicago: Aldine Publishing Co.): 97-140.

Freeman, A. Myrick III, 2003, *The Measurement of Environmental and Resource Values: Theory and Methods*, *Second Edition* (Washington, DC: Resources for the Future).

Hanemann, W. Michael, 1984, "Discrete/Continuous Models of Consumer Demand," *Econometrica* 52 (May): 541-562.

Hanemann, W. Michael and Edward Morey, 1992, "Separability, Partial Demand Systems, and Consumer's Surplus Measures," *Journal of Environmental Economics and Management* 22: 241-258.

Mäler, Karl G., 1974, *Environmental Economics: A Theoretical Inquiry* (Baltimore: John Hopkins University Press for Resources for the Future).

Mansfield, Carol, George Van Houtven, F. Reed Johnson, Zachary Pekar, Jui Chen Yang, Timothy Gibson, and Douglas Crawford Brown, 2002, "Behavioral Reactions to Ozone Alerts: Study Design for Assessing Parents' Values for Protecting Their Children's Health," Presented at the Southern Economic Association Meetings, November 24-26, 2002.

Neill, Jon R., 1988, "Another Theorem on Using Market Demands to Determine Willingness to Pay for Non-traded Goods," *Journal of Environmental Economics and Management* 15 (2): 224-232.

Shogren, Jason F., 2001, "Children and the Environment: Valuing Indirect Effects on a Child's Life Chances," *Contemporary Economic Policy* 19 (October): 382-396.

Smith, V. Kerry and H. Spencer Banzhaf, 2003, "A Diagrammatic Exposition of Weak Complementarity and the Willig Condition," *American Journal of Agricultural Economics* (forthcoming).

Smith, V. Kerry and George Van Houtven, 2003, "Recovering Hicksian Consumer Surplus within a Collective Model: Hausman's Method for the Household," *Environmental and Resource Economics* (forthcoming).

Willig, Robert D., 1978, "Incremental Consumer's Surplus and Hedonic Price Adjustment," *Journal of Economic Theory* 17 (February): 227-53.

ABSTRACT

Family Decision Making and the Value of Preventing Childhood Developmental Impairment

Alan Krupnick[1], Wiktor Adamowicz,[2] Ann Bostrom[3] and Sandra Hoffmann[1]
Resources for the Future, [2]U. of Alberta, [3]Georgia Institute of Technology

Objectives. This multi-disciplinary study has three objectives: (1) to provide more comprehensive valuation of reducing risk of childhood developmental impairment neurotoxins by estimating parental willingness to pay (WTP) to reduce lead-paint hazards; (2) to develop more accurate methods of eliciting parental WTP to protect children's health by testing the impacts of using intra-household resource allocation models rather than unitary household models; and (3) to develop more systematic approaches to designing non-market survey instruments by using mental models research to better understand the underlying decision processes.

Experimental Approach. Current research on parental WTP to protect children's health from environmental hazards models household choice as an action by a single decision-maker with a unitary utility function and a pooled budget. A large body of literature in economics and sociology of the family and in cognitive psychology suggests that these assumptions are flawed, and may lead to incorrect measures of parental WTP to protect children's health from environmental hazards. The study proposed here tests these assumptions.

Environmental neurotoxins can have permanent effects on children's intelligence, motor development, and attention. New research suggests that environmental exposure may be a larger contributor to developmental impairment than previously thought. Lead paint exposure will be used as the representative neurotoxin hazard. Lead paint abatement options provide an excellent vehicle for testing unitary versus bargained household models in estimating parental WTP to protect children's health. EPA currently relies on more limited cost-of-illness and human capital estimates of the benefits of protecting children from lead and other neurotoxin hazards.

This study has two phases. The first phase will elicit "mental maps" of parents' risk perceptions and decision making about reducing their children's health risks from lead paint. Thirty couples will be interviewed individually and as a couple about their risk perceptions, definitions of the decision problem and choice set, and roles in family decision making. Results from phase one will guide development in phase two of a attributed-based/conjoint contingent behavior survey of 250 couples to assess the influence of individual and shared parental risk perceptions and preferences on household choice of lead-paint abatement to protect children.

Expected Results. This research will produce estimates of individual and household willingness to pay for reduced risk of developmental impairment in children. It will also generate measures of the degree to which estimates from bargained and unitary household models differ. Insights into household decision making on environmental health risk issues from the valuation and mental model results will help guide future valuation efforts as well as neurotoxin risk communication programs.

Robin Jenkins's Policy Discussion of
"Benefit Cost Analysis and the Entanglements of Love"
Theodore C. Bergstrom, In progress, October, 2003 and
"Weak Substitution, Environmental Vulnerability, and Choice"
V. Kerry Smith, Mary F. Evans, H. Spencer Banzhaf, and Christine Poulos, October 6, 2003

I will discuss Prof. Bergstrom's paper, the Smith et al paper, and end with connections between the two.

Prof. Bergstrom's paper has two broad objectives. The first is to try to figure out how introducing an important aspect of real life into economic models will affect our interpretation of benefit-cost analysis (BCA). The important aspect is family relationships. How will accounting for the interdependencies of families affect how we conduct and interpret economic analysis?

Prof. Bergstrom's second objective is to simultaneously figure out whose willingness to pay (WTP) is valid when conducting BCA for families. (Should one examine the mother's, the father's or the child's WTP? Some combination thereof?)

Both of these objectives are important to policy-makers. EO 13045, issued by the Clinton administration in the late 1990s, directed policy makers to consider health risks to children. This introduced a new need for policy analysts to try to look separately at the impact of government policies on children. It turns out that this task is virtually impossible if one relies on the traditional tools of economics in which individuals maximize utility for themselves subject to a budget constraint. Children are not rational beings, they don't have access to the family budget and they don't operate as individuals. Children are inherently connected to family. Thus, a new question arose - how do we manipulate the tools of economics to find out whether we've improved the welfare of children? Prof. Bergstrom's paper is taking a stab at the answer. He's really asking how we use the traditional tools of economics (utility theory and BCA) in new, perhaps somewhat unconventional ways, in order to get at valid measures of the effect of policy on child and family welfare.

Prof. Bergstrom starts by reminding us what exactly we get from BCA, anyway. BCA tells us whether a project leads to a potential Pareto improvement (PPI), a situation where the winners could potentially compensate the losers. In the simplest model of an economy, to check for a potential Pareto improvement, one would just examine whether the sum of individuals' WTP exceeds the cost of the project.

Prof. Bergstrom suggests, however, that the relationship between BCA and potential Pareto improvements gets more complicated when you move away from the individual and try to model family relationships in which people care about one another's well-being. To understand the effect of family relationships he shows that we need a theory of household decision-making. He proceeds to examine households of different compositions and with different decision structures:

Household Compositions, Decision-Making and WTP		
Household Composition	**Decision-Making Structure**	**Whose WTP for BCA?**
Single parent	Dictator	Parental WTP
Child-free Couple	One is dictator	Dictator's WTP
	Unanimity	WTPm=WTPf
	Bargaining solution	Minimum WTP
Married couple with child	One is dictator	Dictator's WTP
	Unanimity	WTPm=WTPf
	Bargaining solution	Minimum WTP
Divorced couple with child (independent budget)	Separate dictators	WTPm + WTPf
Extended family (independent budget)	Separate dictators	Follow rule for nuclear family + WTP ext fam

Prof. Bergstrom offers a new criterion for passing BCA when family decisions are being accounted for: **The sum of parents' WTP must exceed project costs to pass the BC test for parents.** As he has worded it, the criterion clearly reveals whose perspective counts: parents. We don't add a child's WTP to parents or we'd end up accepting projects that don't meet this new criterion.

For each household construct, what he considers to get the third column, is a hypothetical government project that improves the health of children for the families with kids, or that improves the health of the female member for the families without kids. As the third column shows, to assess this criterion, depending on the composition and decision-making structure of the family, the policy analyst would collect WTP from different people. The following rule seems to fall out:

Rule for appropriate WTP:
"Interview all families who care about the child and operate with independent budgets. For each of those families,
a) if there is not unanimity between the adults, then count the WTP of the head (dictator) or, if there is not a head, count the WTP that is the lower of the two adults' WTP;
b) if there is unanimity, count the WTP from either individual."

This suggestion has immediate practical implications. For analysts relying on parental data, one must now acknowledge that parental WTP might overestimate if the un-interviewed spouse has a lower WTP, or for divorced couples, it might underestimate since really we should be adding in

the WTP of the ex-husband/wife too.

In addition, if there are segments of the population that are more or less likely to be characterized by one or another of the family compositions or the decision-structures then those sub-populations might be systematically mis-represented. One can see how this is likely to be a real problem since culture is an important determinant of the prevailing power structure within the family. For example, the head-of-household model is probably more common among Hispanic families than non-Hispanic. In addition, certain family compositions are more likely among certain sub-populations – single parent families among inner-city African Americans. Thus the mis-representation of certain sub-cultures might follow from neglecting to model one or another of the possibilities in the table.

Prof. Bergstrom points out that in most of these models, a project could pass the BC test for parents, and children could experience a decline in utility. This is because of substitutability between the child's health and consumption. The parent might reduce the child's consumption by enough to offset an improvement in the child's health, so that the child's utility actually declines. Prof. Bergstrom offers a practical example of an impoverished family that requires a child to go to work once his health is restored. This is a disturbing possibility -- that by relying on parental WTP, policy makers might choose projects that ultimately lead to reductions in the direct utility and consumption of children.

This is a disturbing possibility and one to be aware of but, as Prof. Bergstrom himself points out, at least for western cultures, it seems like a remote one. In response to the child's improved health, the parent would have to reduce the child's consumption by an amount big enough to offset the improved health. I would assert that in Western families, health might be more realistically modeled as complimentary to consumption, since healthy kids probably eat more, are more concerned about social pressures regarding dress and play, and probably recreate and travel more too. In other words, improving children's health is not likely to decrease their direct utility for families in the U.S. at least not as a consequence of parents considering child health to be a substitute for child consumption of goods.

An interesting question is that as this paper develops, will there be more examples in which direct child utility declines even though a project passes the parental BC test -- examples more applicable to western cultures. Are there certain cases or circumstances where we should be concerned about this outcome? Clearly there are outlier parents – drug addicts, for example -- for whom we should be reluctant to rely on parental WTP to represent the well-being of children. But are there also broader cases or situations for which we should exercise caution when relying on parental WTP?

A final comment about the importance of Prof. Bergstrom's work. It goes beyond implications for the interaction of policy and child health. In fact, just as children are inherently family members, so too are most adults. In the final section of his paper, Prof. Bergstrom suggests adding together the WTP of divorced couples and adding to the WTP of a nuclear family, the WTP of extended family members with independent budgets. These suggestions

imply that for adults we should also collect WTP for adult health from family members operating in households with independent budgets. For cultures like ours and most Western ones, where household budgets are fairly independent but there is a great deal of concern for extended family member's health, Prof. Bergstrom's work suggests that individual WTP probably understates society's valuation of health.

Turning to the Smith, Evans, Banzhaf and Poulos paper . . .

The major thrust of this paper is to propose a new framework for revealed preference research. As a policy analyst, after reading the first few paragraphs of this paper, I was excited about the significance of what they suggest. Clearly, analysts would find very useful yet another set of conditions under which we could tease out values of non-market goods or more specifically, for the EPA, environmental goods. The paper sets out to show that when certain conditions hold regarding the interrelationship between a private market good and a substitute environmental good, economists can study the demand for the private good and tease out valuation of the environmental one.

We are accustomed to thinking this way about complimentary goods – for years economists have been learning about the value of natural areas by examining recreation demand; e.g., the demand for fishing or for beach house rentals. As the authors of this paper highlight, Rick Freeman's much-cited text on Measuring Environmental and Resource Values covers the conditions surrounding weak complementarity. His latest version of the text, however, has all but eliminated discussion of weak substitutability. The authors intend, with this paper, to correct this oversight and to encourage economists to study, or at least be more aware of, this neglected category of market products – that is, substitutes for difficult-to-value non-market goods.

For weak complementarity, two conditions must be met. Together they suggest that when the price of the complementary market good is at or above its choke price (so that consumption is zero), changes in the environmental good have no welfare significance. Weak substitutability also requires that two conditions hold. Together they suggest that there exists a level of consumption for the private good (the weak substitute) above which improvements in the non-market environmental service have no value. In other words, if the price of the weak substitute is high enough, consumption of the market good is low and increases in the environmental good are valuable.

One can easily see how positing the conditions for weak substitutability might offend environmentalists. The hypothetical is that there is a market good that is a substitute for some set of environmental services. If the price of that market good drops low enough, then people no longer have a need for the non-market environmental services. Years ago when I taught an undergraduate environmental economics course, following Tom Teitenburg's text book, I would start the semester by exposing students to ideas of the "optimists" and the "pessimists." The citations to these schools of thought go back at least to the 1960s. The pessimists believed the world's population was growing at a rate that would outstrip resources and lead to a collapse

with starvation and misery. The environmentalists in the classroom were sympathetic to these ideas. The optimists believed that advances in technology would outpace resource depletion and eventually replace the need for the natural world. Generally speaking, the optimists viewpoint was disliked by the environmentalists.

Drawing a connection to the paper, in general, studying goods that are complementary to nature, suggests that we have a continuing need for the natural world. These products complement or enhance our experience of the natural world. However, studying goods that substitute for nature, the very notion that goods can completely replace the demand for environmental goods, is in concert with the ideas of the optimists. The weak substitutability condition that there are prices of the market good below which changes in the quality of the environment are no longer valued might be more offensive to some people than the choke price assumption behind weak complementarity.

This leads to a practical question: Is the weak substitutability condition more restrictive or less representative of practical experience as well? I think not as long as one carefully characterizes the case under study. The private market good substitutes for a **specific subset of services** derived from an environmental good, and generally not for the entire set of services. The paper would do well to emphasize this caveat.

This point is illustrated by the examples identified by Smith et al in which weak substitutability might lend insight. Both are drawn from the existing literature and in both, the substitute goods are medical care and certain aspects of a cleaner environment. In other words, the consumer is substituting an averting or mitigating behavior for a cleaner environment, in both examples for their children. Agee and Crocker look at chelation therapy as a substitute for removing lead from a child's environment; Mansfield et al study removing children from the out-of-doors as a substitute for cleaning the air of pollutants that exacerbate asthma.

Certainly, if analysis of weak substitution can be as illuminating for environmental valuation as weak complementarity has, as the authors suggest, then the examination of averting and mitigating behaviors takes on a new importance. A simple application that comes to mind might be to examine purchases of bottled water as the route to valuing cleaner water.

I wanted to draw attention to an important connection between the two papers. In the final section of the Smith et al. paper, the authors highlight that in order to impose weak substitution in analyses of the costs of mitigating or averting behavior, analysts need more information about decision-making within the household . Clearly, Bergstrom's paper is a step in the right direction to fulfilling that need.

To close, both papers advance the state of knowledge regarding non-market valuation which ultimately is useful to policy-makers as an input into BCA; for analysts at the EPA for BCA of environmental policy. The Bergstrom paper gives insight into family decision-making and into the appropriate willingness-to-pay measures for children's health. The Smith, et al paper suggests a new approach to valuing environmental services, and in the examples they

offer, the value stems from protecting human health. These papers make clear contributions to advancing policy-makers understanding of the benefits of human health protection.

Mark Agee's Comments on: "Weak Substitution, Environmental Vulnerability, and Choice," by V. Kerry Smith, Mary F. Evans, H. Spencer Banzhaf, and Christine Poulos.

This paper examines usefulness of identifying a threshold demarcating consumption and non-consumption of substitutes linked to an environmental amenity to assess Hicksian welfare changes associated with amenity improvements.

The paper provides an interesting and seeming applicable framework for use of revealed preference data to valuing environmental amenities. A highly desirable aspect of this framework is its focus on observable goods demands (i.e., does not necessitate specification/estimation of the consumers' underlying home production technologies) and the consumers' observed tradeoffs between these demands.

The first part of the paper lays out a detailed graphical portrayal of the concept of welfare measurement using weak substitution as another form of demand interdependency that arises for private goods and environmental services. I found the approach very interesting, and most of my comments/questions focus on my thoughts as to how the approach might be put in to practice.

Application-related questions/comments:

On page 14, z is referred to (from this point on) as a "mitigating good." Given this characterization, are we saying (or can we say) in effect that the separability notion is one of partitioning the classes of goods (averting and/or mitigating) that identify as substitutes in the individual's choice set—of which are then linked to the individual's utility and environmental amenity through perceived health, or risk to health, etc.? If so, then it seems that a binary choice model might be useful in identifying and estimating key thresholds of interest (you refer to these thresholds as points of discontinuity in the role for q) as well as identify the conditional demand for z?

- e.g., suppose X^a represents a substitute (averting) good such that, at a threshold, z (a mitigating good) is no longer demanded. A binary choice model similar to Agee and Crocker (1996) and Dickie and Gerking (1991) could be used to identify either positive or zero consumption levels of z by estimating:

- $v(P, P_z, P_r, q, \gamma, m)$ + error,

or

- $v[P_z, P, q, \gamma, m_v(P, P_z, P_r, q, \gamma, m)]$ + error,

where $v(\cdot)$ is an econometric specification of a conditional indirect utility function for z, and γ accounts for personal characteristics.

Estimation of the probability of "positive" consumption would also provide

- an estimate of P_a^*,

the price of X^a that reduces z to zero (or the individual's estimated choice index below the threshold), and

- an estimate of (the Marshallian) z as a function of prices, income and personal characteristics: $z(P_z, P, P_r, q, m, \gamma) = -\frac{\partial v(\cdot)/\partial P_z}{\partial v(\cdot)/\partial m}$ based on the analysis of the threshold of no consumption.

The paper also mentions the importance of demographic and physical traits that contribute to characterizing individuals who are differentially vulnerable to the environmental hazard in question. If the model can be estimated as a threshold of zero consumption, it seems that the level of q is also important in identifying key aspects of the price P_a^* since the relationship between goods linked to q can change as q changes:

- e.g., with reference to child lead exposure—there are burdens of body lead such that mitigation is not even considered; burdens such that mitigation together with exposure reduction are absolutely necessary (complementary goods); and burdens (likely close to the threshold) such that exposure reduction can suffice for chelation therapy (substitute goods).

- This of course would change the sign of the coefficient for P_a^* (i.e., only used, used with chelation, used as a substitute to chelation) in the estimated $v(\cdot)$ expression.

In these cases, it seems that a set of interactions between P_a^* and q (and between P_a^* and some of the personal characteristics) would be necessary to identify any critical thresholds for potential sign-changes of the P_a^* coefficient, as well as to correctly identify (calibrate) measures of z_0, z_1, and P_a^* for estimation of WTP.

Another estimation issue is encountered if X is multidimensional. If X is a vector, is it in general necessary to identify the entire vector of X in order to correctly specify z?

- e.g., z might denote (one-dimensional) medical care consumption that mitigates air pollution respiratory ailments; and X denotes exposure reducing activities that, at a point, may render mitigation either very improbable or unnecessary..

Can the problem of multidimensionality of P in z be dealt with adequately by estimating a specification of $v[\cdot]$ (e.g., like the second $v[\cdot]$ specification mentioned above) that includes each individual's observed fraction, m_v, of their budget allocated towards z and a single selected element of P (such as the price of air purifiers)? If so, since m_v is a function of all prices, q, income, and personal characteristics, it seems that

one might want to account for possible correlation between m_v and the error term to avoid potential bias of the P and other coefficients in the binary choice regression (see e.g., Rivers and Vuong (1988) or Wooldridge (2002), p. 472 for cross section data; or Jones and Landwehr (1988) for panel data).

--or--

Might it be possible to construct a one-dimensional averting activities index (like intensity of use from a limited set of "most used" averting activities) with an accompanying averting cost index that condenses P to a single measure for estimation?

Sources:

Agee, Mark, and Thomas Crocker. 1996. "Parental Altruism and Child Lead Exposure: Inferences from the Demand for Chelation Therapy," Journal of Human Resources 31: 677-691.

Dickie, Mark, and Shelby Gerking. 1991. "Willingness to Pay for Ozone Control: Inferences from the Demand for Medical Care," Journal of Environmental Economics and Management 21:1-16.

Jones, J. Morgan and Jane T. Landwehr. 1988. "Removing Heterogeneity Bias from Logit Model Estimation," Marketing Science 7(1): 41-59.

Rivers, D., and Q.H. Vuong. 1988. "Limited Information Estimators and Exogeneity Tests For Simultaneous Probit Models." Journal of Econometrics 39: 347-66.

Wooldridge, Jeffrey M. 2002. Econometric Analysis of Cross Section and Panel Data Cambridge, Mass.: MIT Press.

Suppose X^a represents a substitute [averting] good such that—at a threshold—*z* [a mitigating good] is no longer demanded. A binary choice framework can identify either positive or zero consumption levels of *z* using:

- $\nu(P, P_z, P_r, q, \gamma, m)$ ***+ error***

--or--

- $\nu[P_z, P, q, \gamma, m_v(P, P_z, P_r, q, \gamma, m)]$ ***+ error***

where

$\nu(\cdot)$ **= specification of a conditional indirect utility function for *z***

γ **= personal characteristics**

Estimation of the probability of "positive" consumption provides

- **an estimate of** P_a^*
- **the price of X^a that reduces *z* to zero (the individual's estimated choice index below the threshold)**
- **an estimate of (the Marshallian) *z* :**

$$z(P_z, P, P_r, q, m, \gamma) = -\frac{\partial \nu(\cdot)/\partial P_z}{\partial \nu(\cdot)/\partial m}$$

Mark Agee's Comments on: "Benefit Cost Analysis and Entanglements of Love," by Theodore C. Bergstrom

This paper develops a creative approach to addressing some important questions about aggregation of willingness to pay (WTP) when WTP has altruistic connections between individuals within a family. Four specific questions are addressed:

- When considering a public project that increases children's health, would it make sense to calculate the aggregate value of that project by summing the average parent's WTP by the total number of parents; or
- if parents (within the same household) reveal WTP differences for the project, should the maximum or minimum of these values be used (exclusively) to represent the entire household's value; and
- should children's own values be accounted for and included; and
- should values of family members outside the immediate family be counted.

These questions are addressed within the frame of a specific criterion: is the project Pareto improving? That is:

- is it possible to implement the project and assign project costs to families in such a way that, given the household decision structure, no family member is made worse off and at least one member is made better off?

The initial part of the paper rearticulates this criterion by demonstrating that, if parents have complete control of their family incomes, if one parent is forced to pay more than his/her WTP for a project (implying project costs exceed the sum of parental WTP), then the project is not Pareto improving—and thus the benefit-cost test for parents provides a clear-cut gauge of satisfaction of the Pareto criterion.

The benefit-cost test is then applied to a variety of family structures involving one or two parents with and without a single child, and a variety of family preference and distribution structures common to the household decisionmaking literature. The paper acknowledges that the preference and distribution assumptions of some of the models do not accurately portray modern U.S. households. However, a few of these models had structural assumptions that came (somewhat) close to the frameworks found in the limited number of literature studies examining parental valuations of specific child health attributes. Specifically:

- A parent decisionmaker who allocates after-tax income between own utility-enhancing consumption and child utility-enhancing consumption;
- A parental utility function that derives utility from parental perceptions of child utility; and
- An exogenous index of child health that enhances parent utility through child utility.

Although the basic structure of these examples (within the context of the above assumptions) preclude the possibility of endogeneity of parental decisions regarding child health and child numbers, the paper demonstrates a result similar to prior studies:

- that parental WTP for a child's health improvement equals the health improvement times that adult's MRS between own consumption and child health,

but it also finds that

- the parent *not* charged with income allocation may reveal a different WTP for the same child's health improvement—if this parent's marginal utility of child consumption exceeds her marginal utility her own income allocation,

in which case it is appropriate to

- use only the lowest of the two parents' WTPs to represent the family's WTP for the child's health improvement.

If a general household welfare function is used to link all family members' utilities, then:

- only one parent's WTP (which equals the other parent's and child's WTP) should represent the entire household's maximum WTP for the child health improvement.

However, if family utilities are disconnected (e.g., because of divorce or independent budget arrangements) then

- the sum of adults' WTP for the child health improvement satisfies the benefit-cost test.

These results suggest that aggregate WTP estimates of children's health improvements may inaccurately reflect true WTP if calculated simply as the "average" of parent's WTP multiplied by the number of U.S. parents. The paper highlights the importance of how specification of family allocation processes and the structure of family preferences can potentially impact aggregate WTP estimates. The question of whether to aggregate by U.S. family numbers, number of U.S. parents, or by U.S. parents categorized by their degree of "control" over household resources provides a good start to better understanding how aggregate WTP for children's health improvements can most justifiably be approximated. Listed below are some questions/comments that came to my mind upon several readings of the paper:

First, I'm not entirely convinced of the plausibility of the argument on page 4 regarding strong substitutability between parental income and child health and the potential for children being made worse off by their parents. The argument in the household utility function that seems to be missing here is parental own health, which has an intergenerational link (heritability) to child health. Would healthier parents (who likely have healthier children) necessarily behave in this fashion? The question closely parallels the discussion of parent and child endowments and parental investments in child human capital found in Becker and Tomes (*J. of Labor Economics.*, 1986, 4:S1-S39). Although children with better endowments are much more productive/efficient utilizers of human capital investments, better endowed parents still discount investments in their (better endowed) children at lower rates and invest more resources in them.

Second, does the current preference structure impose that parents treat their children as a "collective child" with an expressed, single willingness to pay for *their* improved health? An important question in my mind is—if we randomly select one child in a family would the parent's willingness to pay for that child's health improvement, multiplied by child numbers be the correct sum? For example, suppose expressions (10) and (11) in the paper incorporate a summation over n (predetermined) family members thus creating a household welfare function that varies also by n. With this (seemingly)

minor change, would the appropriate number for benefit-cost analysts likewise be a single number representative of either parent or a single child—or would summing either parent's response across child numbers overstate aggregate WTP by a factor of n-2?

As an additional thought, the current preference structure has child utility nested in parental utility reflecting the concern each parent has for their child, i.e., the parental marginal utility of child utility. Perhaps child numbers could enter as in Becker, Murphy, and Tamura (*J. of Political Economy*, 1990, p.S12-S37) who assume that, with diminishing marginal utility of children, parental altruism is negatively related to child numbers (they also refer to parental altruism as an intergenerational discount rate applied by parents to the per capita consumption of their children). With this modification, how might diminishing marginal utility of children relate to the parents' determination of WTP for child health improvements, and does it provide any insights as to whether parental WTP is for $1/n$ or some other fraction of child numbers?

Finally, another question might be posed with reference to the role of child health in the determination of family utility. Certainly, the current exogeneity of child health in child utility enables ease of mathematical tractability; however, if child utility is modified such that parental allocations of child consumption goods also influence child health, would endogeneity of child health bring about any changes in your results?

Summary of Q&A Discussion Following Session II

Ronnie Leven (EPA, Region 1) stated that in addition to complications in the theory, there are complications in the cases being used, specifically the lead abatement studies. She went on to assert that "chelation therapy is not all good"and actually leads to a spike in blood lead levels and a correlating spike in brain lead levels. She also noted the spike in exposure after certain types of abatement that release lead paint dust into the air to be inhaled, and she listed encapsulation, replacement, delay, improved maintenance, and temporary removal as alternatives to immediate abatement. She ended by saying that the studies on purchasing bicycle helmets also had some complicating health issues.

Scott Grosse (Centers for Disease Control and Prevention) suggested this extension for Professor Bergstrom: Consider willingness to pay for health as a public good and not just as a private good. Dr. Grosse stated that "if we treat health as a private good, then you can't include the willingness to pay for the extended family members who also have children of their own."

Professor Bergstrom responded by saying that he was thinking about how a particular questionnaire related to a specific project should be framed ("it might be a local community project, it might be a national whole airshed, or it might be a national project"). To him it seems reasonable to ask someone, "How would you feel about an improvement in the health of your child?" or, in fact, "How would you feel about an improvement in the health of *each* of your seven children?" and then on to, "How would you feel about an improvement in the health of your nephew?" Professor Bergstrom said that he would propose evaluating that project then by building up from these micro-answers.

Grosse: "But, it may be that people are not willing to spend "n" times as much to help "n" children."

Professor Bergstrom reiterated that you would ask a sequence of questions–regarding the valuation of the health of one's own children, and then on to a relative's children, and even on to a stranger's children–as the building blocks to constructing a policy.

Ellen Post (Abt Associates, Inc.) told Professor Bergstrom that she was bothered by the fact that assigning the willingness to pay value entirely to the dictator parent in a dictatorial household model (and thereby ignoring the other parent because that person, though possibly willing, has nothing to pay with) basically incorporates *intra*-household politics into the policy assessment. She further stated that the belief "if one person controls the budget then only that person's willingness to pay should be counted" would be true for all similar situations of valuing non-market goods.

Professor Bergstrom replied, "Yes, we'll be a little careful about the "should," of course, but essentially if you believe that your policy will not change the structure of household

decisions, then it seems to me a realistic assessment of policy is the consequences of the policy given the household structure." He also provided the reminder that, at least in theory, this is "a benevolent dictator who buys the flavor of ice cream that his children want, and so on," and clarified that "whether that theory is an accurate theory is yet another question."

Post: "I guess I was just saying that the idea of saying that there are different kinds of families with different intra-family politics which determine people's budgets–that has a lot wider implications than, say, just children's health risks.

Bergstrom: "Indeed, I agree."

Laurie Chestnut (Stratus Consulting, Inc.): stated that she was "alarmed" at the suggestion that we should be seeking to determine various relatives' willingness to pay for a particular child's health. Citing a paper concerning altruism from a few years ago by Jones Lee, she said it was her understanding that "if what people care about is other people's utility, then the optimal allocation is just adding up everyone's willingness to pay for their own utility change, and that takes care of everybody." She further stated that it's only in the case where one cares about another's *consumption* of a specific thing that the willingness to pay issue gets more complicated. She closed by saying, "what we care about is each other's happiness, which I think is the primary model of families caring about each other."

Professor Bergstrom responded that it's a bit tricky with a divorced family, where the money transferred from the non-custodial parent to the custodial parent does not go directly to the child–but, a health benefit does. He stated that if a divorced parent voluntarily transfers money to the custodial parent, then one can presume that they had come to an agreement about where marginal money should go.

Chestnut: "So, health might be a special case in some circumstances . . ."

Professor Bergstrom replied that he had put a lot of thought into the relatives issue and concluded that if you're not willing to give your brother some money and say "go buy your child some medicine," maybe that suggests that you find no extra valuation in that. "On the other hand, it may be that institutions aren't well set up."

Don Kenkel (Cornell University) asked this of Robin Jenkins, referring to her discussant presentation of Ted Bergstrom's paper: "Is the potential Pareto improvement criterion really consistent with the Executive Order?–sort of plucking out children and saying, "This deserves *special* attention."" Is this suggesting that even if the social welfare function *doesn't* pass the sufficiency test, we should still worry about children's welfare?"

Robin Jenkins: "In other words, do we want transfers to children?

Kenkel: "Even if the parents don't want them."

Spencer Banzhaf (Resources for the Future) followed with a comment that he said is actually very similar, and he pointed out that toward the end of Ted Bergstrom's presentation, wasn't it ironic that there was clearly a potential for Pareto improvement, but "the husband was made worse off because of the reallocation in the family?" He concluded that "that irony raises this doubt about the usefulness of potential Pareto criteria when the agency in charge of the policy does not actually have the power to reallocate wealth . . ."

Kerry Smith (North Carolina State University) said that he had a different take: "If we take the way Ted characterized the household model and different views of it and look at what he's doing . . . In environmental economics we have sort of two ways we've approached looking at environmental policy: One way has been to present people with programs, and we've said, "Here's a program to do this–or do something else, on a national level." So the commodity that was actually in preferences was however they conceived of the program, and there were all sorts of questions raised with respect to "Well how should people have preferences about programs?" He's actually proposing to commoditize programs, and so he's basically saying, "We're going to introduce a program that has these kinds of consequences and let's be very, very specific about what those consequences are–it would improve the health of child #1 this much, child #2, the wife, and so forth. . . . But the challenge then becomes converting programs into what commodities would change, so that now we have, if we get very, very specific, with individuals–my nephew's health is going to improve; my niece is going to go down the rathole or something–I then have to think about exactly how is that going to be accomplished with the policy or with the program. So, we don't get around the challenge of connecting policy to commodity outcomes–we simply do a better job of identifying the features and preferences that we can recover from behavior and what we can't recover from behavior, whether it's stated or revealed."

Professor Bergstrom responded that if people are asked how they feel about a policy that they don't understand very well (which is typical), "their answer will be of some interest to politicians, of course, but to welfare economists of less interest I suppose. The point is, it always seems to me a useful thing to map policies to consequences as well as possible."

Glenn Harrison (University of Central Florida) directed his comment at Professor Bergstrom but said it was motivated by Robin's "paternalistic"comment that kids are not rational. He said he was disappointed that Professor Bergstrom didn't address the issue of how kids form preferences, and he lamented that all of the discussion so far in the workshop totally ignored the kids preferences and went straight to the parents. "Now, I understand the simplicity of doing that–I understand the logistic ease of doing that, but it's a pretty sloppy way of thinking about the social welfare function. Just because kids

are vertically challenged, we disenfranchise them." He summarized Robin's viewpoint as being that kids don't have well-formed preferences so we can't rely on them, and then he posed the question: "Shouldn't we be thinking about kids' preferences as state-dependent, broadly defined, where the state is the information they have and at a certain age and with a certain information load, they act more consistently over time? And the question, then, is what implications does that have when you start thinking about *social* willingness to pay as distinct from aggregating up from an individual willingness to pay? And indeed, arguably we're interested in social willingness to pay."

Professor Bergstrom replied that it was "a lovely question" that he was reluctant to tackle at the time due to the upcoming lunch break. He offered to speak to Dr. Harrison on the subject later.

J.R. DeShazo (UCLA) commented, "Ted's presentation, I think, illustrates that when we elicit values from parents what we're really eliciting are household preferences for investments in children's health and the benefits that accrue to children as perceived by the households, so that households can express the schedule of values for the children as a function of the number of children they have, as a function of the governance structure within the household. I simply want to point out that this approach misses the fundamental theoretical construct that we're trying to recover, which are the net benefits to the child of the policy. I think we should just recognize the disconnect between what we can measure and theoretically what we optimally would like to know in order to set the optimal policy."

Professor Bergstrom replied, "That is a technical issue that in a sense I *did* address. The question is to what extent do the parents' evaluations of the child help represent the child's interests. . . If there's complementarity between the child's health and the child's conception, then definitely the parent will *always* be acting in the child's interest, and if there's not, there are cases where indeed the parent will *not*, despite the fact that the parent cares about the child's interests. I can't give you magical answers, you know."

Sandra Hoffman (Resources for the Future) said, "It seems to me that there is a logical inconsistency in the discussion we're having here, because on the one hand we're saying that children do not have well-formed preferences–they haven't developed the judgment that's able really to express preferences–and yet when we look at parents' preferences, we're looking at the child's utility entering the parent's utility function. So there's a transformation going on there. We're not directly measuring the child's benefits, and maybe we don't want to. So it seems to me that J.R.'s point still stands–we're at best getting at a transformation, and perhaps . . ."

Professor Bergstrom clarified that this is not such a difficult issue–the parent's and the child's notions of what's good for the child "needn't be identical–it's as simple as that."

Glenn Harrison said, “But no one here has assumed that the child does not have well-formed preferences.”

Professor Bergstrom responded, “I’ve never made such an assumption. I think it’s a very interesting issue, but again, I don’t want to spend our lunchtime here.”