

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.

¹A 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.

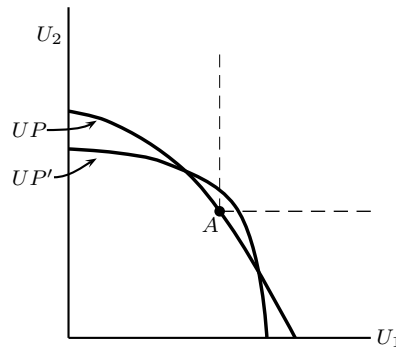
step back and examine the logic of benefit-cost analysis as applied to families in which some individuals care about the well-being of others.

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 a time.

²It 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.

³The 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.

⁴If 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:

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

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

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

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

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

⁵A 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

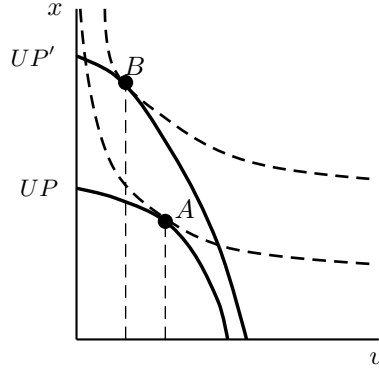
Definition 3 *A parent's preferences are benevolent if the parent's utility function is $U_i(x_i, v_i(k_i, h_i))$ where v_i represents the child's preferences.*

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

⁶Though 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

$$U_A(t) = v_A(x_A(t), h_A(t)) + aU_B(t-1) \quad (4)$$

and Bess's happiness is determined by the function

$$U_B(t) = v_B(x_B(t), h_B(t)) + bU_A(t-1). \quad (5)$$

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:

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

and

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

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

⁷Miles 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

$$U_A(x_A, x_B, h_A, h_B) = v_A(x_A, h_A) + av_B(x_B, h_B) \quad (8)$$

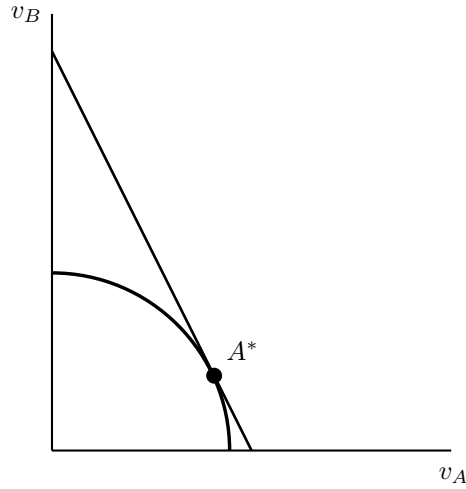
subject to the budget constraint $x_A + x_B = M$, where M is the family income. For given health levels and family income, the possible distributions of the private utilities, v_A and v_B are points lying on or below the curved line in Figure 4.1. Archie’s indifference curves over these possible distributions 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?

⁸In 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 utility⁹ for their own consumption of private goods and that Bess's private utility function is of the special form,

$$v_B(x_B, h_B) = f_B(x_B + kh_B) \quad (9)$$

⁹We 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 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 is 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

¹⁰This 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

$$W(U_A^*(x_A, x_B, h_A, h_B), U_B^*(x_A, x_B, h_A, h_B)). \quad (10)$$

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

$$W(U_A^*, U_B^*) = \alpha U_A^* + (1 - \alpha)U_B^* \quad (11)$$

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 functions 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

¹¹Abram 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.

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

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

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

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 is 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 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

¹²This 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.

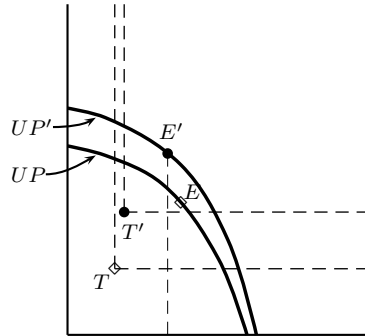
In our example with Archie and Bess, suppose that health levels are determined by public policy while the allocation of private goods is determined by bargaining. Suppose that the utility that each would obtain in the absence of agreement depends on total household income and on the health of each. Where 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

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

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 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

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

¹³Lundberg 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.

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

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

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

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

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

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

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

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.

¹⁴More 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.