Making Interpersonal Comparisons of the Value of Income
With a Hypothetical Auction

Stephen J. Schmidt

Department of Economics, Union College

Revision date: June 15, 2019

1. Introduction and Statement of Claim

If social policies are to be evaluated, at least in part, on their consequences for people, then we need some way to compare the value of gains to some people with the value of gains to others, to determine which options are socially most valuable. In some cases, the choice seems obvious. For example, if we can spend $1,000,000 to operate a medical clinic in a poor neighborhood, or $1,000,000 to operate a horse racing track in a wealthy neighborhood, most people would agree that the medical clinic is the better allocation, in some moral sense of what a "better allocation" means. If we can make that kind of comparison, then we can say what option is best in a policy choice where different options benefit (or harm) different people. Without some method to compare the value of gains and losses to different people, we can only fall back on the Pareto principle's argument that a policy is good if it creates gains for someone and losses for no one, or assert that the merits of a policy are independent of its consequences. The former option allows us to solve very few practical problems, and the latter is implausible.

Unfortunately, most practical problems are less obvious than the preceding example, and so we need some formal way to compare the value of gains and losses. One way to evaluate the gains and losses of policies to different people is to use cost-benefit analysis. To do that, we measure the gains and losses according to amount of income people would give up to get the gains, or would require as compensation to accept the losses, and recommend a policy if the sum of the gains exceeds the sum of the losses. This can be done, and is widely done, but it relies on the implicit assumption that a gain (loss) of one dollar creates the same value (harm) no matter who gets it, at least for the purpose of evaluating the policy. This is implausible in many cases. We have strong intuitions that a dollar produces more value when given to a poor or starving person than it does when given to a wealthy person, or that it produces more value when given to a person with expensive needs (e.g. a terminal disease which is expensive to treat) than to a person who does not have such needs. But cost-benefit analysis would

select the race track for the wealthy over the health clinic for the poor if the wealthy are willing to pay more to get the track than the poor are willing to pay to get the clinic.

It would be very helpful to have a method of establishing the relative value of money to different people; these could be used as weights in cost-benefit analysis, so that it would recommend policies which create net value, rather than those which create net dollar benefits. This would prevent the analysis from recommending policies which create small gains for wealthy people but larger losses for poor people, who cannot afford to pay much money to avoid those losses. One way to do this, widely discussed, is to use a social welfare function that explicitly accounts for the differences in the amount of well-being or welfare that different people get from a marginal dollar of income (Adler 2016a). This requires making interpersonal comparisons of welfare. Much effort, going back at least as far as Harsanyi (1953), has gone into finding a method to establish objectively valid interpersonal comparisons of welfare. This research agenda has been extremely fruitful, but has not produced its ultimate goal; a means of comparing the value of money or goods to different people that is widely accepted.

One could reject the idea that such comparisons should be based on welfare or well-being, and evaluate policies on some other metric of how much value those policies create for different people, such as Sen and Nussbaum's capabilities approach (Sen 1993). But this would not get rid of the problem of comparing the amount of value that different people receive from an increase in income or other policy change. It would only change it from evaluating the well-being of different people, to evaluating whatever alternative consequences we evaluate instead for different people.

Many economists, and some non-economists as well, have concluded that it is impossible to make interpersonal comparisons of the relative value of gains and losses to different people (Robbins 1932, Hausman 1995). But this is clearly too strong a conclusion. It is easy to make subjective comparisons of the value of income to different people; each of us does it each time we choose to donate money to one

3

charity and not another (assuming those donations have no direct effect on our own well-being). But it seems wrong to base collective policy decisions on of one individual's subjective perception of the relative value of gains and losses to different people, particularly if there is a lot of disagreement about it, as there is likely to be (Greaves and Lederman 2018). How could we select the one person whose judgment will be used? Subjective judgments of interpersonal comparison of value therefore do not have a great deal of moral standing; one person's view of the comparison is no better than another's, and so if they disagree as they often do, we do not know which comparison to use.

If we are unwilling to use a subjective method of making interpersonal comparisons, and unable to find an objective one, then it is not clear what we can do. In practice we often rely on using cost-benefit analysis, and either making rough attempts to weight the costs and benefits of policies to disadvantaged groups (often done in Europe) or to make no attempt to weight costs and benefits, but to give distributional issues some kind of separate evaluation when making policy decisions (often done in the United States). In both cases, policy analysts use whatever kind of distributional judgment seems right to them, that will not cause the public or its elected representatives to dismiss their analysis. But this does not provide any objective validity either, and there is not much agreement on how weights should be calculated, or other types of distributive judgements should be made (Adler and Posner 2006, p. 188).

In this paper, I propose a different solution to the problem, which is to find a means of aggregating the subjective valuations of different individuals, on an equal basis, into a group-level valuation that reflects the collective judgment of all members of the group. Such a valuation also does not have any objective validity, as it differs depending on who is included in the group. But it reflects the view of all the members of that group, with each member having equal weight in the decision, rather than the individual view of any one of them. And if the composition of the group changed, or if one member had a change of views about the valuation, then the collective valuation would change, as it should do if every member's opinion should have weight in determining the valuation.

4

I argue that such an aggregate comparison of gains and losses to policy is appropriate when determining policy on behalf of all members of that group. In particular, when a government makes policies that cause gains and losses to the members of its society, it is appropriate to compare those gains and losses using a standard that aggregates, on an equal basis, the valuations held by each of its citizens. Such a standard gives equal moral respect to the views of each person who has agency in the question. In the absence of an independently justifiable standard, a standard that gives each person's subjective view equal standing is preferable to any method that privileges the views of some people over others.

The method of aggregating preferences that I propose is related to Okun's idea of transferring money in leaky buckets (Okun 1975). Suppose an otherwise-disinterested observer prefers to give $1 to person A instead of giving it to person B. That choice could be based on their perception of A's and B's well-being, but could be based on something else, whatever the observer thinks is relevant. That choice implies that the observer regards the marginal value of a dollar given to person A as greater than the marginal value of a dollar given to person B. (Throughout this paper, I will use the term "person" to refer to someone who has an income which has a value relative to the income of other persons, and the term "observer" to refer to someone who has a subjective view of what those values are. For some problems some individuals may be both persons and observers in these senses, but I will keep the two roles separate in the analysis.)

However, if the bucket used to carry the money to person A is "leaky", so that not all the money reaches person A, the observer may face the choice of giving, say, 60 cents to person A (due to a 40 cent leak in the bucket) or $1 to person B. In this case the disinterested observer may prefer giving $1 to person B. If so, the observer must regard the marginal value of a dollar to person B as greater than the marginal value of 60 cents to person A. If the observer chooses consistently, then there should be some amount M for which the observer is indifferent between giving M cents to person A and $1 to person B.

If so, then the observer regards the marginal value of M cents to person A as equal to the marginal value of $1 (or 100 cents) to person B; and this implies that the observer is willing to trade money between the two at a rate of 100/M. For instance, if the observer is indifferent between giving 75 cents to person A and $1 to person B, then the observer weights the value of income to the two of them in a 4 to 3 ratio. More generally, if the transfer amounts to persons A and B are T_A and T_B, and the observer is indifferent to giving income to A or B, then the observer's relative weight for income to person A must be T_B/T_A.

Such choices, if made consistently across different people, will reveal this observer's subjective comparison of value of gains (or losses) of income to all people. For example, suppose there are three people in society, and further suppose that transfers yield 25 cents to person A (a 75 cent leak), 50 cents to person B, and one dollar (no leak) to person C. If under these conditions an observer is indifferent between giving money to any two of these people, then that observer's subjective valuation of income for each of the three people must be in the ratio 4 for person A to 2 for person B to 1 for person C. This valuation of incomes to different people is the one this observer would use, if given the problem of determining whether a policy that gave income to some people and took it from others created a net increase value for society.

To produce a collective valuation, we need to aggregate the subjective valuations of every observer in the group into a single comparison that reflects the views of all the observers in the group. I propose that this can be done by using a hypothetical auction market that establishes a collective value for the income of each person, in the same way that actual markets establish a collective value - a price - for goods that reflects the views of all buyers and sellers in the market about the value of those goods. This hypothetical auction would work in the following manner. An auctioneer announces amounts of income that reach each person when transfers are made to that person, or is taken from them when transfers are made away from that person; in the previous example, 25 cents for person A and 50 cents for person B. Each observer decides whether, given these transfer amounts, they would like to transfer

income between people, and if so, how much. In the example, an observer who felt that the value of income for persons A and B was equal would want to transfer money from person A to person B, taking 25 cents from person A but giving 50 cents to person B, which has more value in this observer's subjective view. But a different observer who felt that income for person A was three times as valuable as income for person B, so that 25 cents of income for person A is worth 75 cents to person B, would want to transfer money from person B to person A, since 25 cents for person A creates more value than 50 cents for person B does, in this observer's subjective view. An observer who values income to person A and B in the ratio 2:1 would not want to transfer money between person A and B in either direction; 25 cents for person A has the same value as 50 cents for that observer. Note that for each observer, there is some set of transfer amounts (determined up to a positive multiple) at which that observer does not want to transfer money between person A and B, and that this ratio is the inverse of the ratio of their value of income for the two people.

Once the auctioneer has announced particular transfer amounts for each person in the society, every observer decides how much income to transfer between all of the persons. Depending on their subjective views about the value of money to each person, some observers would transfer money in one direction and others in another direction. However, I show that there will always be some set of transfer amounts at which the net demand for transfers of income would be zero for every person. I argue that, just as the transfer amounts at which a single observer would not want to make any transfers imply the subjective interpersonal comparisons of that observer, the transfer amounts at which the group collectively does not want to make any transfers should be taken to imply the collective valuation of the group of observers for income for each person.

The resulting valuation is, I claim, the appropriate one to use when spending money that is collectively and equally controlled by the observers. Since no objective valuation exists, and each observer's subjective valuation has the same moral weight as every other observer's, a valuation that

7

gives each observer equal weight in determining the valuation used is the morally correct way to proceed. This aggregate valuation can be calculated using only observable choices about how each observer would allocate income between persons, so it is not conceptually difficult to actually compute the implied values of income to different people and use them as weights in cost-benefit analysis. The claim that these are the right weights to use does not depend on any moral assumption about the value of income to different people except that each observer's subjective value should be given equal weight in determining it. In that sense, it respects the moral agency of each observer whose money is at stake, rather than impose ideas about the value of income about which the observers might disagree. It delegates the question of value to the members of the decision-making group in much the same way that markets delegate the question of the value of goods to buyers and sellers of those goods.

The rest of this paper provides the answers to four questions. First, under what conditions will observers have consistent subjective valuations of money to different people in society that we could observe? Second, how can a hypothetical auction produce a collective value for the group by aggregating the values of the individual observers? Third, why would this collective value be a reasonable one to use in comparing the values of gains and losses to those individuals in public policy decisions? Fourth, how could this valuation be calculated in practice? Sections 2 through 5 discuss each of this questions in turn; section 6 concludes.

2. Subjective Comparisons of the Value of Income to Persons

We cannot determine a collective valuation of income for different persons for use in policy choices unless we have subjective valuations of individual observers to aggregate. To measure those subjective valuations, we can observe the choices an observer would make when being allowed to choose between giving different amounts of money to those persons. As long as the observer makes those choices

8

consistently, the choices will imply a ranking of possible distributions of income among persons, which will imply relative values for income for each person.

More formally, let there be a set of N persons, called a society, who have characteristics x_i and incomes y_i. A social state is a set of characteristics and incomes for all persons. Let there be M observers, called a group, who are being asked to choose how to allocate money to different persons; that is, to choose among the social states. If an observer's choices are complete, reflexive, and transitive, then the observer will have a preference ranking over the set of all possible social states. This ranking will be different for different observers, if they have different views about the value of income to different people in the society. Let those preferences be represented by a value function $U_j(y)$, different for different observers, where y is a vector containing the values of y_i for each person in society.

Note that this value function represents the preferences of observers over income distributions. It has no necessary relationship to the preferences of the persons themselves about how to spend that income, nor to their well-being; it is just the preferences of the observers about the distribution of income among the persons. In particular, these preferences are not the same as Harsanyi's concept of extended preferences. They are not based on comparing the well-being of different persons, none of whom is the observer. We are not asking the observer whether s/he would prefer to be person A with characteristics x_A and income y_A or person B with characteristics x_B and income y_B. Therefore, they do not require the observer to understand what it would be like to be persons A or B with their preferences and characteristics, which is a problem with extended preference theories (Greaves and Lederman 2018). They do not even require the observer to think about the well-being of person A and B; they only ask the observer whether s/he thinks it would be better, in the observer's view, for person A or person B to have money. The observer might use his or her evaluation, possibly inaccurate, of the well-being of the two persons, but need not do so. Any method of ranking the social states that the observer may wish to use is validSince we are not arguing that these preferences have any kind of objective validity,

9

they only need to reflect whatever information this observer believes is significant in making the evaluation.[1]

We will, however, need to impose two conditions on these preferences in order for them to have some kind of moral meaning. First, we need them to be informed preferences, in the sense that the observer has given due consideration to the choice and taken all relevant factors into account. In particular, whatever information any observer regards as significant has to be included in x.[2] Second, we need them to satisfy anonymity; the preferences must be indifferent between two social states if they give the same combinations of characteristics and incomes to different persons. That ensures that the observer cannot favor one person over another; either herself if she happens to be one of the persons in the society, or family or friends if they are persons in the society.

We can then ask each observer to solve a relatively simple problem; given the total income in society Y_T which is just the sum of the given y_i, and the ability to transfer money costlessly between persons, find the distribution of income that maximizes her subjective value. That is, maximize $U(y^*)$ subject to the constraint that $\Sigma y_i^* = Y_T$, where y^* is the vector of the observer's optimal incomes for each person. The first order condition for the problem is that the marginal value of additional income to each person will be equal, given her characteristics. This will normally require the observer to transfer some income away from persons who have relatively low marginal value for income (in the observer's subjective view) and give it to persons who have relatively high marginal value for it, until the observer thinks the marginal value of income to each person in society is equal.

---

[1] Someone who accepted the extended preferences view, but not Harsanyi's claim that extended preferences should be the same for all individuals, could use a similar procedure to aggregate those extended preferences.
[2] If we want to impose further limitations on preferences, we can do that by limiting variables that can be in X. For example, if we wish to follow Roemer (1998)'s suggestion that Tiny Tim's sunny disposition should not affect his ability to receive money to fund a wheelchair, we can require that sunniness of disposition not be an X variable. If so, then observers must place the same value on income for all individuals with the same X characteristics, regardless of their sunny dispositions, or lack thereof. This can also be used to eliminate offensive preferences, e.g. those where the recipient's race affects the marginal value of income to that person.

This is not, however, the problem we need to solve; we need to determine the observer's relative value of income for different people at the level of incomes they start with. To find that out, we can ask each observer to solve a different, somewhat more complicated problem. In this problem, money cannot be moved costlessly; only a fraction of the income being sent to a person is received by that person, due to a leaky bucket. Let there be a transfer amount t_i for each person; if the observer gives $1 to person A, then person A only receives t_A; the rest is lost to a leak, the amount of which is 1 minus the transfer amount. The leak amount is effectively a price for transferring money to person A, and the transfer amount will function very much like a price below. Conversely, if the observer gives $1 less to person B, then person B only loses t_B, because the leak amount is no longer lost.

In this problem, observer j maximizes $U_j(y^*)$ given the constraint that $\Sigma\ t_i\ y_i^* = \Sigma\ t_i\ y_i$. The previous problem is just a special case of this problem in which $t_i=1$ for all persons. (Or t_i could be any constant value, but we can normalize that value to 1. More generally, rescaling all transfer amounts by a constant does not change the problem or its solution.) The solution to this problem is a set of transfer demand functions $y^*=f(t)$. The first order conditions for this problem are of the form $MU_{ij}(y^*) = \lambda^*t_i$ where MU_{ij} is observer j's value of marginal income for person i, and λ is the Lagrange multiplier of the constrained optimization problem; the shadow value of one unit of additional income (before any leaky transfers) for society. From these conditions it follows that if $t_i = MU_{ij}(y)$ for each person – the transfer amount for each person is proportional to the marginal value of additional income for that person at the starting incomes – then the starting incomes will solve the first-order conditions, and so the observer will not want to transfer any income to or from any person. Then the relative value of income for each person, according to this observer's view, is just $1/t_i$. We can infer this valuation by observing the choices the observer makes among different possible income distributions.

3. Aggregating Subjective Valuations into Collective Valuation

Of course, that valuation will be different for different observers. The next step in the problem is to aggregate those different valuations into a single valuation that represent the collective judgment of all the observers about the relative value of income to different persons in society. We could average them, or do something else equally simple, but there is no particular justification for doing so. Or we could have the different observers vote on particular valuations and see if one can command majority support. But this sort of aggregation will run into Arrovian impossibility problems; it is unlikely that one valuation will be preferred by a majority of observers over all others, as the conditions that could guarantee a majority favorite (e.g. a one-dimensional preference of alternatives and single-peaked preferences, which would assure a median observer's preferences would always command a majority) are not likely to hold. So we need to take another approach.

The approach I propose is to sum up the individual demands of each observer, to produce the aggregate demand for transfers of income at a given set of transfer amounts. There will exist a set of transfer amounts for which the aggregate demand for transfers is equal to zero for all persons in the society. With these transfer amounts, the group of observers as a whole does not want to transfer any income between members of society. Some observers will want to transfer incomes to particular persons, others will want to transfer income away from those same persons, based on their different subjective interpersonal comparisons of the value of income to persons. But when each observer is allowed to propose the transfer amount that seems best to him or her, those proposals exactly offset one another. Just as the transfer amounts that cause an individual observer to desire no transfers indicate that observer's value for income of each person, so these transfer amounts, at which the group desires no transfers in aggregate, indicate the group's value of income of each person. They reflect the aggregate judgment of all the observers in the group about the value of income to different persons. Their reciprocals are then the appropriate weights to use in comparing the relative value of gains and

losses to the persons in society, when the collective view of the group of observers is the one that should be used to make that comparison.

We can show that such a set of transfer amounts must exist by using a fixed point theorem, which is structurally similar to the theorem that proves the existence of a set of prices that clear all markets in an exchange economy.

Theorem 1: Let there be a group of M observers, and a distribution of income among N persons. Each of the observers has a value function $U_j(y, x)$ where y is a vector of incomes for each of the N persons and x are characteristics of each person, and j indexes observers. Each observer selects a preferred $y_{ij}*$ that maximizes $U_j(y*, x)$ subject to the constraint $\Sigma t_i y_{ij}* = \Sigma t_i y_i$ where i indexes persons. Let $Z_{ij}(t) = y_{jj}* - y_i$ be the demand for transfers of income to each person (from each person, if $Z_{ij}(t) < 0$). Let $Z_i(t) = \Sigma Z_{ij}(t)$ be the aggregate demand for transfers to person i. Then there exists a value of t such that $Z_i(t) = 0$ for each person.

The proof is found in the appendix.

Since the group does not wish to transfer income between two people when the relative cost of doing so is t_A*/t_B*, it follows that $1/t_A*$ and $1/t_B*$ are appropriate measures of the collective relative value of income for persons A and B, in the sense defined above.

4. Why Collective Valuation? Why This Valuation?

I have argued so far that individual observers can form consistent valuations of income for different people, and I have provided a means for calculating a collective valuation that aggregates, equally, the valuations of all individuals whose opinions have moral standing. Why should we use that particular valuation to make actual comparisons of gains and losses to persons from policy choices? I offer three

arguments. First, the aggregate valuation gives each subjective observer's personal valuation equal weight, which is a requirement of respecting the moral agency of each person. Second, there exists no objectively correct valuation that can be used to show that the aggregate valuation is wrong. Third, the aggregate valuation has several properties that we would like our social valuation to have.

Suppose that M=1; that is, the money being distributed is clearly owned by a single observer prior to being distributed. Then surely that observer is entitled to use his or her own subjective valuations, rather than be required to use anyone else's. Our respect for his or her moral agency requires that we accept the observer's decisions based on his or her view of the value of income to different people, when the decision will only affect the allocation of money that is properly controlled solely by that observer.

By extension, when the money in question is under the collective control of the group of observers with each one having an equal right to determine how the money is used, then we should use some kind of collective valuation of the incomes of people affected by the decision. This case might arise because all observers in the group are citizens of a government, and the problem is to decide how to spend money that is owned by that government. Or it might happen because the government is setting taxes, and everyone agrees that the income of private persons is justly subject to being claimed by the government through the tax mechanism (within any other limits we might want to impose on the tax mechanism, of course). In that case, the collective valuation via the hypothetical auction is a reasonably natural way to construct a valuation that gives equal weight to the views of each observer.

This is not the only possible way to give each person equal control over the money being distributed. For instance, the observers might divide their collective money equally and let each observer do as he or she sees fit with his or her share. However, that requires each observer to yield his or her partial ownership of an (N-1)/N share of the collectively owned money, in order to receive full control over a

mere 1/N share of it. I do not see any argument we could make that would compel each observer to do this. Rather, each owner should be allowed to keep his or her share of control over all the money, consistent with a like share being available to all the other owners. By extension from this, when government money is being spent, the money should be spent according to a collective valuation over which each citizen of that government has an equal influence. The hypothetical auction is one means of doing that.

One might object to the collective valuation – or to the subjective valuation of a single observer, or to any other valuation - by saying that these collective values of income to different people in society are not the "right" values to use. The observer, or observers if more than one, may not make the right choices about who should get money. A claim of that sort would have to be supported by an objective argument about what made some choices right and other choices wrong. How might such an argument be made?

One fairly natural possibility would be some kind of utilitarian argument based on the well-being of the recipients. But this depends on a means of making interpersonal comparisons of well-being, which is exactly the problem we're trying to solve. If we already had the "right" interpersonal comparisons of well-being, we would not need to calculate an aggregate valuation. We might have objectively true comparisons if there was some objective list of what constitutes well-being, with every item on the list delivering the same amount of well-being to each person, and the value of money being derived from the items on the list one can obtain one with money (unless money is an item on the list, but that seems unlikely; money contributes to well being instrumentally, but in most cases not directly). However, that would rely on a defensible claim about what items belong on the list, and there is no consensus about that. Instead, we might object comparisons of well-being with some kind of mental-state argument about what well-being is, if we could measure the necessary mental states in a comparable way. But that approach to the problem is subject to experience-machine types of problems (Nozick 1974).

Objectively correct comparisons might instead be found using an extended preferences approach, following Harsanyi (1955). This approach assumes that well-being constitutes preference satisfaction, and that we have a way to compare the extent of different people's preference satisfaction. To do that, we have to assume that each observer's valuation if income for different people was based on the observer's evaluation of what it would be like to be person A or B with different amounts of money. But those evaluations might differ from observer to observer. Harsanyi claimed that an ideal observer would have preferences that would be objectively valid, but it is unclear if this ideal observer could exist, or how any real person could know what that ideal observer's preferences were if the ideal observer did exist. There may be a "view from nowhere," in the terms of Nagel (1986), but we cannot know what that view would look like. Adler (2012) argues in favor of using extended preferences where all observers agree about comparisons, and accepting incomplete social preferences where the observers do not agree. Greaves and Lederman (2018), however, argue that is likely to lead to massive incomparability of social states which will make the method unsuitable for solving many real problems.

One could use other methods to produce interpersonal comparisons of welfare, such as equivalent incomes or happiness metrics, both of which are discussed in Fleurbaey and Abi-Rafeh (2016). However, equivalent incomes depends on a choice of a reference price vector (and if non-market goods matter, reference quantities for those goods) to allow comparisons, and it's unclear how such reference prices should be chosen. Happiness metrics rely on the assumption that subjective well-being reports are comparable across different persons, which may not be the case (Fleurbaey and Blanchet 2013).

And all of these approaches are welfarist; they assume that well-being is the only relevant information to use in determining the relevant value of income for different people. This is disputed as well. Capabilities may be more important than well-being, or primary social goods may be what matters, or the decision might not be made on consequentialist grounds at all. On the whole, it seems difficult to find an objective ground from which to criticize the idiosyncratic choices made by each observer about

the value of income to different people. The most we would be able to do is say that our own views

about how the value of income to different people are not the same. If different observers have

different preferences about income distribution that lead to them having different values of income for

different persons, and we cannot establish some of those as correct and the rest as incorrect, how can

we decide whose subjective preferences to use for any particular problem that has to be solved?

If we cannot solve that problem, then we need to find some means of aggregating everyone's

subjective valuation into a collective valuation. There might be many ways of do so; why use the method

of a hypothetical auction to do it? I do not claim that a hypothetical market is the only acceptable way to

do it, but I claim that it is a reasonable way to do it. Real markets are effective ways to produce

collective valuations (market prices) for objects about which different observers (buyers and sellers)

have different subjective valuations. There are many cases where we presently use real markets to

assign values to objects for public finance purposes. For instance, if we want to collect a property tax

that is a percentage of the property's value, we need to agree on a common value of the property to use

in setting the tax bill. Different people have different views about what a piece of property like a house

is worth – whose view should determine the size of the tax bill? In practice, we get a value for the house

(at least conceptually) by taking its value in a property market in which any member of the society can

buy or sell property. As long as the market is reasonably competitive, so that no one individual has a

disproportionate ability to influence the valuations, there is little controversy in using market value as

the basis for the tax – or more specifically, using assessed values that are intended to mimic the value a

property would have if it was sold on the market, which is necessary since individual houses are sold

relatively infrequently. But the market price is the target the assessed value is aimed at. This is a morally

acceptable way to set taxes because it relies on a means of establishing a value – a competitive market –

in which each person in society has an equal ability to influence the market price.

Conventional economic analysis also relies on this logic, when it asserts that (in the absence of externalities) the price someone pays for an object reflects the social value that object's consumption produces. When the market price of an apple is $1 and the price of a banana is 50 cents, welfare economics holds the apple creates twice as much social value as the banana does. Not all individuals agree with that valuation; each consumer has his own personal evaluation of the relative values of the two. People who buy only apples will think an apple has greater value, in their view, than two bananas, and people who buy only bananas will think an apple has less value, in their view than two bananas, and people who buy both will think that an apple is exactly as valuable as two bananas. But we are comfortable asserting that the social value of an apple is exactly twice as great as a banana, and claiming that society should produce at a point where the marginal cost of producing another apple is exactly twice the marginal cost of producing another banana.

The idea of using hypothetical markets to aggregate subjective values to make collective decisions about the social value of non-tradable goods has precedent in the philosophical literature. Most notably, Dworkin (1980) famously argues for using markets to establish the relative value of different resources when trying to distribute resources equally among individuals. Dworkin wishes to ensure that all members of society have equal bundles of resources, but he cannot give each one an identical bundle because some of the resources (talents and the absence of handicaps) are not transferrable between people. Thus, he needs some means to establish the value of each resource, so that the transferrable resources can be transferred in a way that gives every member of society a bundle of resources with the same total value. But different people place different values on resources, depending on their preferences, and no objective valuation of the resources is available. Dworkin argues, "the idea of an economic market, as a device for setting prices for a vast variety of goods and services, must be at the center of any attractive theoretical development of equality of resources." Dworkin's particular idea is a hypothetical auction in which people bid for resources using clamshells, which have no value except to

serve as currency in this auction, and each person has the same number of clamshells, so that each

person ends up with a bundle he (weakly) prefers to the bundles of everyone else. Because each person

has the same number of clamshells, each participant is an equal in Dworkin's market; the hypothetical

market is not subject to the problem that all real markets face, that people with greater wealth can get

more in the market. The prices that clear this market are, Dworkin claims, the morally relevant values

for each of the resources. Dworkin's market has been criticized for the way in which it holds people

responsible for preferences they may not have chosen freely (for instance, Roemer 1985, Cohen 1989,

Arneso 1989) but this criticism doesn't touch the basic idea of markets as aggregators of individual

values for resources.

The problem of establishing the value of income to different people is functionally similar. The goal is

to maximize the value of the income distribution rather than to equalize allocations; but the problem is

essentially the same because in order to do that, we need a way of placing a value on the income of

each person. Different observers have different views about the value of income to each person, and a

hypothetical market can produce a collective value for income to each person in the same way one

produces a collective value for resources in Dworkin's problem. Just as giving each person an equal

number of clamshells ensures that the collective values of resources reflect each individual's subjective

value of those resources equally, so giving each person the same ability to redistribute income in my

hypothetical market ensures that each observer has an equal share of agency in establishing the value of

income to different people.

Further, the hypothetical market produces valuations of income for different purposes that satisfy

several properties we might want a collective valuation to have.

1. Unanimity; if every observer agrees about the relative value of income to different persons, then the

    collective valuation will produce the same relative values. This is because at the transfer amounts at

which every observer wants to make no transfers, the aggregate demand for transfers will necessarily be zero.

2. Pareto: if all observers think that the income of person A has a higher value that the income of person B, then the collective valuation will value person A's income more highly. If it were not so, then person B would have a lower transfer amount than person A, and each observer would want to transfer money from B to A. But if so, that transfer amount doesn't have zero net transfers between A and B. As a special case, this implies that if the observers all have preferences with diminishing marginal value of income, and person A and person B are identical except for their incomes, then the collective valuation will place a greater weight on income for the one with less income.

3. Strong monotonicity; if one observer changes his view so that they believe the value of income to person A is higher (lower), then the collective valuation will give a strictly higher (lower) value to person A also. This is because at the original valuation, there was zero net demand for transfers to person A. When one observer raises his value of income to person A, that observer will want to transfer more money to person A. That will result in a positive net flow of transfers to person A; person A will have to be given a lower transfer amount to restore zero net transfers, and that will increase the collective valuation of A's income. (With a large number of observers, the effect of changing the subjective valuation of any one of them will be very small.)

4. Non-dictatorship; the collective valuation is not simply the valuation of one observer. This follows directly from strong monotonicity; if any one observer changes her subjective valuation, the collective valuation changes.

5. Anonymity; the collective valuation is independent of the identities of the observers. This is because permuting the individual valuations of the observers will not change the net demand for transfers; they all count the same way in the summation of transfer amounts.

6. Transitivity; it will consistently rank the value of income to different people. If it finds income for person A more valuable than income for person B, and income for person B more valuable than income for person C, then it will find income for person A more valuable than income for person C. This is because it places a numerical value on income for every person, and those values will obey transitivity.

There are two properties we might like it to have that it does not have. First, it does not satisfy a condition like independence of irrelevant alternatives; the collective valuation's relative value of the incomes of persons A and B depends on more than each observer's subjective relative value of the incomes of persons A and B. This is not surprising; if it did satisfy a form of IIA then a version of the Arrow impossibility theorem would apply. Greaves and Lederman (2017) have proposed that authors working in the extended preferences framework should consider methods which violate IIA in order to avoid Arrow's theorem in that framework. While my framework is not an extended preferences framework and so their advice is not strictly relevant, the same basic logic applies; there is not nearly as much reason to want IIA to hold as there is to want conditions like Pareto optimality and non-dictatorship to hold. Similarly, Hammond and Fleurbaey (2004) suggest that,

> faced with diversity of ethical opinion, dictatorship of ethical values appears inevitable if one is going to have a (complete) social welfare ordering satisfying some form of Pareto criterion and of independence. But in this field as well, relaxing independence might be a promising way ahead, though as yet it has been little explored.

and this approach is going down that road. (But for a contrasting view, see Adler 2016b).

Second, it is strategically manipulable. An observer who realizes that the collective valuation will be different from her individual valuation can misrepresent her preferences, bidding for higher transfers to people whose income she values more than other observers than she really wants, in order to raise the net demand for transfers to those people and hence reduce the transfer amount that will produce zero net transfers. This is also not surprising since the intuition of Gibbard's theorem (1973) applies. It seems unlikely that any non-dictatorial collective valuation exists that is not strategically manipulable. Any system for attempting to calculate the collective valuation of a group of real observers will have to provide an external incentive for observers to report their preferences correctly. This ought not to be very difficult; we will return to this point below.

5. Calculating Collective Values of Income in Practice

In order to use these collective marginal valuations in actual cost-benefit calculations, we need to be able to compute values for them. Calculating exact values is practically impossible; it would require knowing the preferences of all citizens of the government making the decision. However, if we knew the preferences of a representative sample of citizens, we could calculate approximate weights based on that sample. Although they would not be exactly the right weights to use, they would be substantially better than not weighting at all, or weighting using cruder methods. They would also be empirically derived weights that would not depend on the subjective judgement of the person doing the weighting, beyond the decision to use this method of calculating weights rather than another. In that sense they have the same moral appeal as market-based prices do – given the decision to base prices on individual preferences, the calculation of the prices themselves is objective.

It is not conceptually difficult to elicit an individual's preferences for giving income to different people. All we have to is ask that individual to make a choice between two possible social states; would s/he prefer to give $10 to person A or to person B. As long as the individual chooses consistently, the answers to such questions will reveal the individual's preferences. It is not difficult to ask individuals to make such choices, either in a face-to-face interview or perhaps over the Web. Subjects could asked to make a series of hypothetical choices about giving out income. They should not, of course, be asked about giving income to people they know; so that their preferences will satisfy anonymity, they should have to choose based only on information about the X characteristics and incomes of people who would receive income. For example, they might have to choose whether to give $50 to Ms. Smith, a single mother with two children and an annual income of $25,000, or $75 to Mr. Jones, a married person with an annual income of $60,000 and a medical condition that requires treatment that costs $5000 per year. To encourage subjects to report truthfully, one might carry out one of the subject's choices by giving real money to a person with the characteristics described in the choice. Such surveys would not be more complicated than similar surveys which aim to provide national-level information about subjective well-being (Benjamin, Heffetz, Kimball, and Szembrot 2014) or other surveys aiming to get information "beyond GDP" about national economic performance.

With data on choices made by individuals in that way, we could estimate a parameterized value function representing individual preferences. These preferences might depend on both the characteristics of the people receiving money and the characteristics of the individuals expressing the preferences; age, ethnicity, political affiliation, etc. Using such an estimated value function, we could then calculate the marginal valuations of income that would clear the hypothetical market, and use those to calculate weights to use in cost-benefit analysis.

The following detailed example shows how this could be done. Consider a society with six people (or six equally-sized groups of people) who have incomes, and four observers (or four equally-sized groups of similar observers) who value the people's incomes. Each observer's preferences for each person's income depends on how much income that person has, with decreasing marginal returns, and a characteristic of that person X which takes the value 0 or 1. People with X=1 get greater value from additional income than people with X=0; it might represent a health condition that requires treatment, or anything else that means one person gets a higher value from additional income than another. The observer's preferences also depend on a characteristic of the observer Z, which takes the value 0 or 1. Observers with Z=1 believe the marginal value of income is higher for all people than observers with Z=0 do; perhaps they place more weight in material conditions of life in evaluating social states. After observing many choices by the observers about how income should be distributed, we believe that the observer's preferences can be represented by a function of the form:

$$V_{ij} = \beta_0 + \beta_1 {}^* Y_j + \beta_2 {}^* Y_j^2 + \beta_3 {}^* Y_j {}^* X_j + \beta_4 {}^* Y_j {}^* Z_i$$

where V_{ij} is the ith observer's value of income for person j. [3]The marginal value of additional income for person j, in the opinion of observer i, is thus

$$MV_{ij} = \beta_1 + 2 {}^* \beta_2 {}^* Y_j + \beta_3 {}^* X_j + \beta_4 {}^* Z_i$$

Suppose observed choices show that observers 1 and 2 have different preferences, and hence different values of the β parameters, than observers 3 and 4 have. Specifically, let their preferences be such that

$$MV_{1j} = 5 - 0.05 {}^* Y_j + 2 {}^* X_j + Z_1$$

$$MV_{3j} = 5 - 0.02 {}^* Y_j + X_j + Z_3$$

---

[3] I use a quadratic term for mathematical convenience only, so that the marginal value of income is linear in the income. In practice it would probably be better to use an Atkinson-type function instead: U(y) =

and similarly for observers 2 and 4. Observers 1 and 2 believe in stronger diminishing returns for income than do observers 3 and 4, and they also place a greater value on X. Further, Z=1 for observers 2 and 4 and Z=0 for observers 1 and 3, so that all four observers have different preferences, with one difference controlled by Z and one by differences in β parameter values.

Suppose that the six people in society have incomes of 20, 20, 40, 40, 60, and 60 respectively (think of them as annual income in thousands, if you like) and for persons 1, 3, and 5, X=0, and for persons 2, 4, and 6, X=1. If income can transferred costlessly, then all observers will want to transfer income from persons with high income to low income and from persons with X=0 to persons with X=1. Table 1 shows the amount of income that each observer would want to transfer to each person in society if transfers were costless. Observers 1 and 2 want to transfer the same amounts because Z makes only a linear difference in the marginal value, and similarly for observers 3 and 4, but observers 1 and 2 want to transfer different amounts than observers 3 and 4 do. In particular, with their preferences, observers 3 and 4 want to transfer $5 from person 1 to person 6, while observers 1 and 2 with different preferences do not want to transfer any money to either person 1 or 6. If transfers were costless, then the four observers, both individually and collectively, would want to transfer the most money to person 2 (who has the lowest income and X=1) and the most away from person 5 (highest income and X=0). Collectively they would want to transfer 10 income to person 6 and away from person 1.

If transfers are leaky, then there are transfer amounts at which each observer would not make any transfers. However, the amounts are different for each of the four observers since their preferences are different. Table 2 shows the transfer amounts at which each of the four observers would not want to make any transfers of income between the six people in the society. Those costs differ even between observers 1 and 2, and between observers 3 and 4, because the ratio of each person's marginal value of income is different for them. This implies that the weights that each observer places on the income of the six people in society are different for different observers. All four observers agree that person 5 has

25

the lowest relative value of income, and the transfer amount on his income is normalized to 1. The first observer believes that it is equally good to give 50 cents to person 1 as to give $1 to person 5 (because the ratio of marginal utilities is 2 to 1). This observer is willing to accept a 50-cent leak in the transfer to person 1 - and so gives a weight of 2 to income for person 1. However, the second observer, with different preferences, believes it is equally good to give 60 cents to person 1 as to give $1 to person 5. This observer is only willing to accept a leak of 40 cents to transfer money to person 1, because the ratio of marginal utilities is 5 to 3 rather than 4 to 2. This observer's subjective weight for the income of person 1 is only 1.667. Similarly, the third observer places a weight of only 1.211 on income for person 1 and the fourth observer places a weight of only 1.167 on income for person 1. Each observer has a subjective value of the income of each person in society, but they are not the same because each observer's preferences about income distributions are different.

However, Theorem 1 shows that there exists some set of transfer amounts – leaks from the buckets going to each person – at which the sum of the desired transfers by the four observers to each of the six people will be zero. These can be calculated numerically. Table 3 shows the set of transfer amounts, applied to each of the four observers, at which each observer wants to make some transfers between people, but the sum of those transfers is equal to zero for every person in the society. The cost of transferring money to person 1 is 23.2 cents – that is the size of the leak in his bucket. At that cost, observer 1 wants to transfer $0.85 away from person 1, observer 2 wants to transfer $0.08 away from person 1, observer 3 wants to transfer $0.50 away from person 1, but observer 4 wants to transfer $1.43 to him. The sum of these four transfers is zero. Similarly, the sum of the transfers to each of the other persons is 0 when this set of transfer amounts is applied. At these costs, the observers collectively do not want to transfer any money between any person in society (even though individually each of them do). The weight of each person's income, calculated by the reciprocal of the leak in that person's bucket, thus represents the observers' collective judgement of the value of income to each person. Table 4

shows the subjective valuations of income by all four observers and the collective valuation as implied

by the transfer amounts in table 3. All four observers agree that person 5, who has an income of 60 and

X=0, has the lowest relative value of income, and so that person ends up with the lowest value

collectively, normalized to 1. All observers agree that person 2, with an income of 20 and X=1, has the

highest value for income. Person 2 gets a collective weight of 1.642, which is intermediate between the

weights that the four observers have individually (ranging from 1.375 for observer 4, who sees the least

difference in the relative weights of income to different people in the society, to 3.000 for observer 1,

who sees the most difference).

My claim is that when these four observers collectively and equally control money that can be

distributed among these six people, these weights are the correct ones to use in determining how

income should be distributed. When projects are being done by a government in which those four

observers are equal citizens, then these weights are the correct ones to determine whether a project

has positive or negative net benefits for the society.

Making such calculations for a real society would be substantially more complicated. The calculations

would have to include more characteristics of people that that affect observer's valuation of their

income, and more characteristics of observers that affect their valuations. There would have to be

research to determine which characteristics have a significant enough effect on valuations to be worth

including in the model. There would be more groups of people and more groups of observers, and they

would be different sizes so they would to be weighted by population in calculating net transfers.

However, all of these problems are technical difficulties rather than conceptual ones. With a certain

amount of survey research and a certain amount of calculation – not more than is currently done to

calculate price indices or unemployment rates – this approach could produce estimates of collective

valuations of income that could be used in applied cost-benefit analysis. Doing so would produce

decisions about value that would be much more defensible than either using weights selected by individual analysts or using unweighted analysis and incorporating distributional analysis separately.

6. Conclusion

Society needs some way to value gains and losses to different people in order to make policy choices. When those gains and losses have income equivalents, we can do that if we can calculate the relative value of money to different people. Individuals can do this subjectively, but there does not appear to be any viable objective way of establishing values. However, we can establish collective values by aggregating the individual subjective values of each individual. These collective values are not objective either, but they reflect the views of all individuals whose views are included in the aggregation, on an equal basis. These are the right values to use when making social decisions that are the collective responsibility of those individuals.

We can establish individual subjective values by asking individuals how they want to allocate money between different persons at varying relative transfer amounts. The transfer amounts which cause a person to want to make no transfers are the reciprocal of that person's value for each income. We can then establish collective values in a similar manner, by finding transfer amounts for which aggregate transfers are zero. Given those transfer amounts, some observers may want to transfer income from any given person A to B, and others may want to transfer from B to A, but in total, the net transfers are zero. The transfer amounts at which the group collectively doesn't want to transfer income imply the collective value of the group for income to each person, in just the same way the transfer amounts at which each observer doesn't want to make transfers implies that observer's value. The weight that should be placed on each person's income when making social decisions is then just the reciprocal of that transfer amount.

Under reasonable conditions, such a set of weights will exist. We cannot calculate the desired weights exactly, but we can get a good approximation of them by asking a sample of observers about choices they would make, estimating a model of preferences for transfers from the observations, and then calculating market-clearing transfer amounts. We can use the resulting set of weights for cost-benefit analysis that depend on no subjective decision except the decision to aggregate the preferences of observers, whatever they may be, in this particular way. This method produces a collective valuation in which each observer's preferences get equal weight in establishing the collective valuation, and satisfies a number of other conditions that a collective valuation ought to have. Using these weights to perform cost-benefit analysis would give cost-benefit analysis a defensible quantitative means of accounting for distributive concerns, which it currently lacks, and that would help us make more ethical policy choices.

References

Adler, Matthew D. (2016a) "Benefit-Cost Analysis and Distributional Weights: An Overview." *Review of Environmental Ethics and Policy*, 10(2):264-285.

Adler, Matthew D. (2016b) "Aggregating Moral Preferences." *Economics and Philosophy*, 32:283-321.

Adler, Matthew D., and Eric A. Posner (2006). *New Foundations of Cost-Benefit Analysis*. Cambridge MA: Harvard University Press.

Arneson, Richard (1989) "Equality and Equal Opportunity for Welfare." *Philosophical Studies*, 56:77-93.

Benjamin, Daniel J., Ori Heffetz, Miles S. Kimball, and Nichole Szembrot (2014). "Beyond Happiness and Satisfaction: Toward Well-Being Indices Based on Stated Preference." *American Economic Review*, 104(9):2698–2735

Cohen, G. A. (1989) "On the Currency of Egalitarian Justice." *Ethics*, 99:906-44.

Dworkin, Ronald (1980) "What is Equality? Part 2: Equality of Resources." *Philosophy and Public Affairs*, 10(4):283-345.

Fleurbaey, Marc, and Rossi Abi-Rafeh (2016) "The Use of Distributional Weights in Benefit-Cost Analysis: Insights from Welfare Economics." *Review of Environmental Ethics and Policy*, 10(2):286-307.

Fleurbaey, Marc, and D. Blanchet (2013) *Beyond GDP: Measuring Welfare and Assessing Sustainability.* Oxford: Oxford University Press.

Gibbard , Allan (1973) "Manipulation of Voting Schemes: A General Result." Econometrica, 41(4):587-601.

Greaves, Hilary, and Harvey Lederman (2017). "Aggregating extended preferences." *Philosophical Studies*, 174:1163-1190.

Greaves, Hilary, and Harvey Lederman (2018). "Extended preferences and interpersonal comparisons of well-being." *Philosophy and Phenomenological Research*, 96(3): 636-667.

Hammond, Peter J., and Marc Fleurbaey. "Interpersonally Comparable Utility." In Hammond, Peter J., Salvador Barberà, and Christian Seidl, editors, *Handbook of Utility Theory*, vol. 2. Kluwer Academic.

Harsanyi, J.C. (1953) "Cardinal utility in welfare economics and in the theory of risk-taking." *Journal of Political Economy*, 61:434-435

Harsanyi, J. C. (1955) "Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Utility." *Journal of Political Economy*, 63:309–321.

Hausman, Daniel M. (1995) "The Impossibility of Interpersonal Utility Comparisons." *Mind*, 104:473–490.

Nagel, Thomas (1986) *The View From Nowhere*. New York: Oxford University Press.

Nozick, Robert (1974). *Anarchy, State, and Utopia*. New York: Basic Books.

Okun, Arthur (1975). *Equality and Efficiency: The Big Tradeoff*. Washington, D.C.: Brookings Institution.

Robbins, L. (1932). *An Essay on the Nature and Significance of Economic Science*. London: Macmillan.

Roemer, John (1985). "Equality of Talent." *Economics and Philosophy*, 1:151-88.

Sen, Amartya (1993) "Capability and Well-being," in Nussbaum, Martha and Amartya Sen (eds.), *The Quality of Life*. Oxford: Clarendon Press, pp. 30–53

Appendix. Proof of Theorem 1

This proof is similar to the one given in Varian (1984) for the existence of a Walrasian equilibrium for an economy with competitive markets (p. 195-96).

Theorem 1: Let there be a group of M observers, and a distribution of income among N persons. Each of the observers has a value function $U_j(y, x)$ where y is a vector of incomes for each of the N persons and x are characteristics of each person, and j indexes observers. Each observer selects a preferred y_{ij}^* that maximizes $U_j(y^*, x)$ subject to the constraint $\Sigma\ t_i\ y_{ij}^* = \Sigma\ t_i\ y_i$ where i indexes persons. Let $Z_{ij}(t) = y_{ij}^* - y_i$ be the demand for transfers of income to each person (from each person, if $Z_{ij}(t)<0$). Let $Z_i(t) = \Sigma_j\ Z_{ij}(t)$ be the aggregate demand for transfers to person i. Then there exists a value of t such that $Z_i(t)=0$ for each person.

Proof: Normalize all the t's such that $\Sigma_i\ t_i = 1$. Consider the function

$$g_i(t) = (t_i + \min(Z_i(t),0)) / (1 + \Sigma_i(\min(Z_i(t),0))) \tag{1}$$

The economic interpretation of this function is; suppose the Z_i are not all zero. For persons who have net transfers onto them ($Z_i>0$) produce new transfer amounts g_i by raising t somewhat, and for persons who do not ($Z_i<=0$) set g_i equal to t_i. Then renormalize so that $\Sigma_i\ g_i = 1$.

The map g(t) is continuous, and with the normalization we have $0 \le t_i \le 1$ and $0 \le g_i \le 1$, so it maps from a compact subset of R^N into the same compact subset. Therefore it has a fixed point where $g_i(t) = t_i$ for all i. Let that fixed point be t^*. We show that $Z_i(t^*)=0$ for each person at t^*, hence is the desired value.

Given the definition of $g_i(t)$, at the fixed point t^*,

$$t_i^* = (t_i^* + \min(Z_i(t^*),0)) / (1 + \Sigma_i\ (\min(Z_i(t^*),0))) \tag{2}$$

Multiply the denominator across to get

$$t_i^* * (1 + \Sigma_i\ (\min(Z_i(t^*),0))) = t_i^* + \min(Z_i(t^*),0) \tag{3}$$

Subtract t_i^* from both sides and multiply both sides by $Z_i(t^*)$ to get

$$Z_i(t^*) * t_i^* * (\Sigma_i\ (\min(Z_i(t^*),0)) = \min(Z_i(t^*)^2, 0) \tag{4}$$

Sum over i to get

$$\Sigma_i\ [Z_i(t^*) * t_i^*] * (\Sigma\ (\min(Z_i(t^*),0)) = \Sigma_i\ [\min(Z_i(t^*)^2, 0)] \tag{5}$$

The first order conditions for the choice of transfers for each observer say that $Z_{ij}(t^*) * t_i^* = 0$ for every j. Summing this over j, since $Z_i(t) = \Sigma_j\ Z_{ij}(t)$, we have

$$Z_i(t^*) * t_i^* = 0 \tag{6}$$

for all i. Summing this over i gives

$$\sum_i [Z_i(t^*) * t_i^*] = 0 \tag{7}$$

Therefore, the left-hand side of equation (5) must be zero, and so must the right-hand side:

$$\sum [\min(Z_i(t^*)^2, 0)] = 0 \tag{8}$$

This can be true only if $Z_i(t^*)=0$ for all i. Therefore t^* is the value of t at which $Z_i(t)=0$ for each person, hence aggregate demand for transfers to all people is zero. QED.

Table 1: Desired transfers among six people by four individuals, when transfers are costless

			Observer 1				Observer 2				Observer 3				Observer 4			
Person	Income	X	MV	ΔInc	Inc'	MV'	MU	Dinc	Inc'	MU'	MU	Dinc	Inc'	MU'	MU	Dinc	Inc'	MU'
1	20	0	4	0	20	4	5	0	20	5	4.6	-5	15	4.7	5.6	-5	15	5.7
2	20	1	6	40	60	4	7	40	60	5	5.6	45	65	4.7	6.6	45	65	5.7
3	40	0	3	-20	20	4	4	-20	20	5	4.2	-25	15	4.7	5.2	-25	15	5.7
4	40	1	5	20	60	4	6	20	60	5	5.2	25	65	4.7	6.2	25	65	5.7
5	60	0	2	-40	20	4	3	-40	20	5	3.8	-45	15	4.7	4.8	-45	15	5.7
6	60	1	4	0	60	4	5	0	60	5	4.8	5	65	4.7	5.8	5	65	5.7

MV = marginal value of income at starting incomes; ΔInc = desired transfer (negative amounts are transfers of income away);
Inc' = post-transfer income; MV' = marginal value after transfer

Table 2: Transfer amounts that induce each observer to demand no transfers among people

			Observer 1				Observer 2				Observer 3				Observer 4			
Person	Income	X	MV	Leak	Trans	Wt	MU	Leak	Trans	Wt	MU	Leak	Trans	Wt	MU	Leak	Trans	Wt
1	20	0	4	0.5	0.5	2	5	0.4	0.6	1.667	4.6	0.174	0.826	1.211	5.6	0.143	0.857	1.167
2	20	1	6	0.667	0.333	3	7	0.571	0.429	2.333	5.6	0.321	0.679	1.474	6.6	0.273	0.727	1.375
3	40	0	3	0.333	0.667	1.5	4	0.25	0.75	1.333	4.2	0.095	0.905	1.105	5.2	0.077	0.923	1.083
4	40	1	5	0.6	0.4	2.5	6	0.5	0.5	2	5.2	0.269	0.731	1.368	6.2	0.226	0.774	1.292
5	60	0	2	0	1	1	3	0	1	1	3.8	0	1	1	4.8	0	1	1
6	60	1	4	0.5	0.5	2	5	0.4	0.6	1.667	4.8	0.208	0.792	1.263	5.8	0.172	0.828	1.208

MV = marginal value of income; Leak = loss of money in transfer; Trans = amount of money transferred; Wt = implied weight on income of this person for CBA purposes; Trans = 1 – Leak; Wt = 1/Trans.

Table 3: Transfer amounts that produce zero net demand for transfers for each person

			Observer 1				Observer 2				Observer 3				Observer 4				Sum
Person	Income	X	MU	Trans	ΔInc	Inc'	MU	Trans	ΔInc	Inc'	MU	Trans	ΔInc	Inc'	MU	Trans	ΔInc	Inc'	ΔInc
1	20	0	4	0.768	-0.85	19.15	5	0.768	-0.08	19.92	4.6	0.768	-0.5	19.5	5.6	0.768	1.432	21.43	0
2	20	1	6	0.609	18.05	38.05	7	0.609	13.81	33.81	5.6	0.609	-10.6	9.375	6.6	0.609	-21.2	-1.24	0
3	40	0	3	0.869	-11.5	28.52	4	0.869	-8.48	31.52	4.2	0.869	6.228	46.23	5.2	0.869	13.73	53.73	0
4	40	1	5	0.671	7.429	47.43	6	0.671	5.413	45.41	5.2	0.671	-3.9	36.1	6.2	0.671	-8.94	31.06	0
5	60	0	2	1	-22.1	37.9	3	1	-16.9	43.13	3.8	1	12.95	72.95	4.8	1	26.03	86.03	0
6	60	1	4	0.746	-3.2	56.8	5	0.746	-2.98	57.02	4.8	0.746	2.823	62.82	5.8	0.746	3.357	63.36	0

MV = marginal value of income; Trans = transfer amount; ΔInc = desired transfer (negative amounts are transfers of income away); Inc' = post-transfer income desired by observer i.

Table 4: Subjective weights of each individual for incomes to each person, and collective weight

			Weight of observer				Collective
Person	Income	X	1	2	3	4	weight
1	20	0	2	1.667	1.211	1.167	1.302
2	20	1	3	2.333	1.474	1.375	1.642
3	40	0	1.5	1.333	1.105	1.083	1.151
4	40	1	2.5	2	1.368	1.292	1.491
5	60	0	1	1	1	1	1
6	60	1	2	1.667	1.263	1.208	1.340

The weight for each observer is the one that makes that observer's demand for transfers of income for all people equal to zero; the collective weight is the one that makes the total demand for transfers of income to each person equal to zero.