



CPB Netherlands Bureau for Economic
Policy Analysis

CPB Discussion Paper | 364

Distributionally Weighted Cost- Benefit Analysis: From Theory to Practice

Thomas van der Pol
Frits Bos
Gerbert Romijn

Distributionally Weighted Cost-Benefit Analysis: From Theory to Practice

Thomas van der Pol^{1,2}, Frits Bos^{1*}, Gerbert Romijn¹

* Corresponding author, Email address: f.bos@cpb.nl, Tel.: +31 625063936

¹ CPB Netherlands Bureau for Economic Policy Analysis, The Hague, the Netherlands

² Global Climate Forum, Berlin, Germany

Abstract

In CBA practices around the world, benefits are valued regardless of to whom they accrue. This disregards basic economic principles, like declining marginal utility of income, or inequality aversion. This paper argues that if redistribution matters, net benefits must be aggregated using a distributionally weighted CBA. We introduce the building blocks to do so, i.e. a marginal welfare weight selection, a weight normalization and benefit accounting choices, and an analysis of redistribution effects. A case study about child care benefits in the Netherlands is presented as an illustration. We conclude that guidelines on welfare weights, normalization and practical issues are needed to facilitate experiments with case study applications of marginal welfare weights. We also argue that although distributionally weighted CBA more closely adheres to theoretical foundations, it may be demanding, prone to misuse and might distract attention from the decision problem at hand. Ultimately, welfare weighting could undermine political support for CBA. We end with a summary of other ways to highlight income redistribution in CBAs, like an overview of purchasing power effects. If these are insufficient for informed decision-making, distributionally weighted CBA could be considered as a complementary tool for policies that involve major equity effects.

Keywords: Cost-Benefit Analysis – Welfare weights – Inequality aversion – Income redistribution – Child Care Benefits – Normative Economics

JEL classification: D61; D63; H43; H50

1. Introduction¹

In cost-benefit analysis (CBA) practices in countries around the world, benefits and costs are added up irrespective of who gains or loses. Traditionally, the Hicks-Kaldor compensation principle provides a theoretical underpinning for this practice (Harberger, 1978). Hicks-Kaldor compensation implies that the net gains from a project or policy measure are sufficient to potentially compensate the losers without making anyone worse off. This corresponds to a Potential Pareto Improvement. If compensation is costless and a CBA has a positive outcome, i.e.: a positive Net Present Value (NPV) or a benefit-cost ratio exceeding one, then there are net welfare gains from a public project or policy measure for society as a whole.

However, actual compensation may be costly or absent. This provides a general motivation to analyze redistribution effects of a public project or policy measure in a CBA. Compensation is costly if it introduces market distortions. For example, compensation could reduce long-run labor productivity, or reduce labor market supply by a reduction in hours worked or through the participation decision. The social costs of such labor market effects are studied in a separate analysis. Afterwards, these costs may be compared to the changes in the income distribution.

Investigation of ex ante redistribution effects of a public project or policy measure in a CBA may be informative to decision-makers. More specifically, an analysis of income redistribution effects in CBAs could be useful to account for inequality aversion and decreasing marginal utility of income. If decreasing marginal social welfare weights are applied, net benefits accruing to higher income classes are weighted less than net benefits accruing to lower income classes. This type of CBA has been called a distributionally weighted CBA (Boardman et al., 2006).

This paper contributes to the existing literature by examining practical problems that arise when a distributionally weighted CBA is carried out. Whereas recent literature contributions are concerned with the semi-empirical inference of marginal welfare weights that could be used for such analyses (Lockwood and Weinzierl, 2016; Zoutman et al., 2016 and Jacobs et al., 2017) or apply different sets of weights in alternative metrics to value project benefits (Kind et al., 2016), we focus on a number of practical problems that arise if one applies distributional weighting in CBA practice. To this end,

¹ An earlier version of this paper was presented at the 9th Conference of the Society for Benefit-Cost Analysis, March 15-17, 2017, Washington DC. The authors would like to thank Clemens Kool, Joseph Cordes and the members of the Dutch CBA-working group on the Marginal Cost of Public Funds, in particular the two experts Bas Jacobs and Erik Schokkaert, for their comments on drafts of this paper. However, the views expressed are those of the authors and do not necessarily reflect those of these commentators.

we select marginal welfare weights, including those obtained from revealed preferences in the tax system, and we normalize and apply these to a case study on Dutch child care benefits.

The structure of this paper is as follows. Section 2 introduces distributionally weighted CBA. It also provides a brief background on CBA practices, including those related to the marginal cost of public funds. Section 3 elaborates on a marginal welfare weight selection, a normalization choice, practical problems to infer welfare changes across income classes or percentiles, and some benefit accounting issues. Section 4 presents a case study application. Section 5 concludes and discusses implications of the results for the usefulness of distributionally weighted CBA.

2. Background

This section provides a brief and non-exhaustive overview of the theory of distributionally weighted CBA and related CBA practices.

2.1 Theory: individual utility, social welfare and inequality aversion

In this section we investigate the relation between individual utility, social welfare and inequality aversion. From this we arrive at a number of implications that follow from the use of the distributionally unweighted Hicks-Kaldor principle that is commonly used for CBA. These implications are the starting point for thinking about what it means to add distributional aspects to CBA practice (e.g. Cowell and Gardiner, 1999).

First consider individual welfare, or individual utility. Individuals derive utility U from a generalized concept of income Y_i that comprises not only purchasing power but includes wider benefits (and costs) arising from non-market goods and services such as health, the environment, social contacts, safety, general well-being et cetera. For ease of exposition, assume that individual utility has a CRRA form²:

$$U(Y_i) = \frac{Y_i^{1-\varepsilon}}{1-\varepsilon} \quad . \quad (1)$$

Thus, marginal utility of generalized income decreases with its level if $\varepsilon > 0$.

² CRRA stands for constant relative risk aversion. The parameter ε performs a number of tasks. First it describes the degree of concavity of the utility function, i.e. the rate at which marginal utility declines as (generalized) income rises: the income elasticity of utility equals $1-\varepsilon$. This aspect of the utility function is relevant for distributionally weighting in CBAs. At the same time ε describes individual's risk aversion (the rate at which individuals discount uncertain outcomes relative to certain outcomes) and an individual's intertemporal substitution elasticity (the rate at which an individual is willing to forgo current consumption for future consumption). These last two roles of ε are not relevant to the present discussion.

Social welfare aggregates individual utilities and can thus be stated as a function of individual utilities. Again for ease of exposition we assume a CRRA form, i.e.

$$W = \frac{1}{1-\mu} \sum_{i=1}^N (U(Y_i))^{1-\mu} \quad (2)$$

In Eqn. (2) μ reflects society's inequality aversion with respect to individual utility levels. If μ is zero society does not care about the distribution of individual utilities. For CBA this implies that it does not matter whether a project adds utility to poor individuals or to rich individuals. As μ increases more and more weight is added to utility gains of poor individuals at the expense of richer individuals. As μ approaches infinity, the Rawlesian maximin case emerges that uniquely places weight on the poorest in terms of utility.

However, it is hard to observe individual utility levels (or the changes therein), so eqn (2) will not be very helpful to decide whether a policy proposal increases social welfare. Fortunately, eqn (1) can be substituted into eqn (2) to yield

$$W = \frac{1}{1-\rho} \sum_{i=1}^N Y_i^{1-\rho}, \quad (3)$$

with $(1-\rho) = (1-\mu)(1-\epsilon)$. Differencing eqn (3) yields

$$dW = \sum_{i=1}^N \frac{\partial W}{\partial Y_i} dY_i = \sum_{i=1}^N Y_i^{-\rho} dY_i \equiv \sum_{i=1}^N \omega_i dY_i \quad (4)$$

Eqn (4) can be usefully employed to infer the effects of policy measures on social welfare as it involves measuring and monetizing effects on broad income and adding them together using welfare weights ω_i .

A Hicks-Kaldor CBA uses $\rho = 0$ and thus $\omega_i = 1$ for all i (individuals). This derives from its premise as a *potential Pareto improvement*. If individuals that lose out from a particular policy in terms of broad income, can theoretically be compensated by the individuals that gain in terms of broad income and there is still gain left, the same is true for the utility based social welfare measure of eqn (2). So a positive Hicks-Kaldor CBA translates into a positive social welfare gain.

However, the necessary transfers are not usually forthcoming nor can they be costlessly implemented³ (e.g. Working group Marginal Cost of Public Funds and CBA, 2016 and Bos et al. 2017). So potential income or welfare gains may involve redistribution in practice. If we want to account for this in CBA outcomes,⁴ distributionally differentiated weights must be applied.

As a start, it should be noticed that when it comes to assessing the monetary value of the social welfare gain in terms of weighting of utilities, this implies that the Hicks-Kaldor CBA uses $(1-\mu) = 1/(1-\varepsilon)$. In other words Hicks-Kaldor CBA implies welfare weights for individual utilities that are the inverse of the marginal utility of income.⁵ Layard et al. (2008) report estimates of ε based on surveys for OECD countries. However, in such stated preferences studies ε may reflect inequality aversion as well. In the sequel, we will consider:

$$\frac{dW}{dY_i} = \omega_i = Y_i^{-\rho} . \quad (5)$$

This reflects the combination of decreasing marginal utility of income and society's dislike of inequality. Both combine into welfare weights that reflect how the premise of potential Pareto improvement that is implicit in the Hicks-Kaldor approach, is to be qualified if and when distributional issues enter the CBA sums.

2.2 Weights and normalized weights

The weights from Eqn. (5) cannot be directly applied in a distributionally weighted CBA as the weights are relative. We can multiply eqn (2) or (3) by an arbitrary number and they will still be true. This requires some sort of normalization, for example with respect to some reference income (Y_r). Examples of Y_r that are used explicitly by analysts or implicitly by decision makers include the lowest, median, modal or average income. For the reference income, the welfare weight is set to be equal to one, i.e.: $\omega_r = 1$. More generally, normalized welfare weights ω_i^n follow from (Dahlby 2008; Florio 2014):

$$\omega_i^n = \left(\frac{Y_i}{Y_r} \right)^{-\rho} \quad (6)$$

³ Although it must be said that major reforms of the tax system or the public health system in the Netherlands have been accompanied by secondary measures to minimize effects on purchasing power for a rather detailed delineation of the Dutch electorate.

⁴ Which is a big if, see e.g. Working group Marginal Cost of Public Funds and CBA, 2016.

⁵ An alternative interpretation of the Hicks-Kaldor approach to CBA involves the incremental character of a policy intervention within an otherwise optimized system. The envelope theorem then states that the marginal benefits of redistribution fall away against the marginal costs.

In the literature, the average income of the population is sometimes chosen as reference income (e.g. Kind 2016). For this case, the weight for the average income is set equal to one. Suppose that we distinguish ten income groups of equal size (deciles). And suppose that the fourth decile income level is the average. Then obviously Hicks Kaldor weights ($\rho=0$) are one for all income groups (and sum to ten). With $\rho>0$ weights are larger than one for deciles 1-3 and less than one for deciles 5-10. For decile four the weight is exactly one. The average normalized weight over all deciles can be higher or lower than one. This amounts to a different scaling of the average dollar gain in comparison to the unweighted Hicks-Kaldor case. To eliminate this, the welfare weights can also be normalized, such that the weighted average of the normalized marginal welfare weights equals one (Lockwood and Weinzierl 2016). A discrete implementation with J income classes is:

$$\theta(\rho) \sum_{j=1}^J P_j \omega_j = 1 \quad \leftrightarrow \quad \theta(\rho) = \frac{1}{\sum_{j=1}^J P_j \omega_j} \quad , \quad (7)$$

where P_j is the share of people in income group j , ω_j is the non-normalized welfare weight of income group j and $\theta(\rho)$ is a constant for a given degree of inequality aversion. Hence, we may also write

$$\omega_i^n = \theta(\rho) Y_i^{-\rho} \quad . \quad (8)$$

This normalization has the advantage that it reproduces the outcomes of a standard CBA if the net benefits are distributed equally across the population. For other normalizations, this is generally not the case. For example, if the lowest income class is chosen as reference class, standard CBA outcomes are only reproduced if all benefits and costs accrue the lowest income class.

2.3 CBA Practice

Standard textbooks describe ways to highlight redistribution effects in CBAs. For example, a description of distributionally weighted CBA for both income and other distributive effects is included in Boardman et al. (2006) and Florio (2014). The topic of redistribution and marginal welfare weights is also covered in Dahlby (2008) on the marginal cost of public funds.

The benefits and costs that accrue to individuals, their households, income classes or percentiles may be shown in a CBA, as well as the various other effects of a public project or policy on actors, stakeholders or regions. A distributionally weighted CBA goes one step further than this, and assigns different weights to these different groups in order to aggregate distributive effects in the NPV. The mathematics of this procedure is straightforward and is identical no matter which groups are distinguished; see Boardman et al. (2006).

In a CBA, the effects of a public project or policy on the income distribution can first be shown without weighting these effects. To this end, the benefits and costs in a CBA are disentangled in groups or percentiles, and benefits and costs that accrue to people with different incomes are reported. In this way, the effects of a public project or policy on the income distribution are displayed.

Various distributive effects other than income redistribution effects can be reported in a CBA in the same way, such as the distribution of the benefits and costs between different actors, for example between producers and consumers, or the distributive effects between different stakeholders, for example between smokers and non-smokers of a public health policy. Similarly, spatial redistribution between affected regions of, for example, place-based policies can be highlighted as well.

However, despite its apparent simplicity, only few CBA guidelines consider distributionally weighted CBA for CBA practice. For example, the UK Guidelines on cost-benefit analysis ('Green Book') states that: "Any distributional effects identified should be explicitly stated and quantified as far as possible. (...) Where it is considered necessary and practical, this might involve explicitly recognising distributional effects within a project's NPV" (HM Treasury, 2011). The UK Green Book also illustrates how this could be done. Unresolved practical issues remain to perform distributionally weighted CBA. In the Netherlands this has been investigated for child care benefits.⁶

3. Steps to perform a distributionally weighted CBA

We distinguish a number of topics that need to be addressed to perform a distributionally weighted CBA. Section 3.1 describes the selection of marginal welfare weights for CBA. Section 3.2 describes normalization of these weights. Section 3.3 introduces an analysis of income redistribution effects, and some accounting issues that may change the outcomes of a distributionally weighted CBA.

3.1 Selection of marginal welfare weights

There is an abundant literature on marginal welfare weights and their estimates. Marginal utility elasticity estimates are based on stated preferences from surveys, behavioral evidence from life

⁶ This was part of the research carried out by the authors on behalf of the Dutch Working Group on the role of Marginal Cost of Public Funds in CBA (see also Bos et al, 2017 and Jacobs, 2016). In the course of the work a number of simulations were carried out to investigate the effects of differentiated welfare weights. For this we used a microsimulation model for purchasing power, labour cost, social security and income tax (Romijn et al, 2008), and inverse optimal welfare weights (Zoutman et al., 2016). This paper shares the insights gained in the practical problems that may arise when applying marginal welfare weights in realistic settings.

time consumption or other models, and revealed social preferences, such as through government spending or the progressiveness in tax rates. According to the UK Guidelines on cost-benefit analysis (UK Treasury, 2011, p. 98), assuming an elasticity of marginal utility of consumption of 1 is defensible and in line with estimates by e.g. Pearce and Ulph (1995)⁷ and Cowell and Gardiner (1999)⁸. But Evans (2005) argues that this assumption is too conservative, as cross-country income tax rates data indicate that the utility elasticity with respect to consumption is approximately 1.4 for 20 OECD countries. Layard et al. (2008) presents a survey-based income elasticity bandwidth of 1.19-1.34 for 50 countries, including European Union countries. These estimates seem primarily involved with the estimation of ϵ (see eqn 1), but the inclusion of social preferences and the progressiveness of the tax rates suggests that elements of μ may also be involved (see eqn 2).

Recent research attention has focused on inverse optimum methods (Lockwood and Weinzierl 2016). These methods use optimal taxation theory to infer marginal welfare weights from the revealed progressiveness in tax systems. Whereas these semi-empirical estimates may be considered appealing for distributionally weighted CBA, they also suffer from various anomalies, such as a median-voter bias, i.e.: non-monotonicity, and negative welfare weights (Zoutman et al. 2016).

⁷ They estimate a range from 0.7 to 1.5.

⁸ They provide an estimate being just below or just above 1.




Figure 1 Four sets of welfare weights after weighted-average normalization: low and high inequality aversion (left panel) and inverse optimum welfare weights, including smoothed inverse optimal welfare weights (right panel)

In this paper, we select four sets of marginal welfare weights.⁹ Note that we refer to sets of marginal welfare weights for ease of exposition. Obviously, a set of marginal welfare weights may also be a continuous function. The benchmark set contains weights of one, which is in line with the practice of adding up benefits and costs independent of who gains or who loses, as based on the Hicks-Kaldor

⁹ For real-world applications, too many sets would be unworkable and a guideline that specifies these sets is needed for practitioners.

compensation or a dollar is a dollar-principle. This case is represented by the benchmark of no inequality aversion ($\rho = 0$) in Figure 1 (left panel).

The right panel of Figure 1 displays marginal welfare weights of the inverse optimal method based on Zoutman et al. (2016) data. These non-monotonic welfare weights show a number of anomalies. A possible solution to this problem is to fit a marginal welfare function, for example Eqn. (6), to these optimal inverse weights.¹⁰ A least squares regression gives the smoothened marginal welfare weights displayed in the right panel ($\rho = 0.54$). The left panel also shows a 50% lower ($\rho = 0.27$) and a twice as high inequality aversion ($\rho = 1.08$). The latter is broadly in line with the recommendation in the UK Guidelines on cost-benefit analysis (HM Treasury, 2011, p. 98), but lower than the estimates by Evans (2005) and Layard et al. (2008).

3.2 Normalization

Marginal welfare weights have to be normalized to perform a distributionally weighted CBA. As an example, consider an imaginary society that consists of a low (Y_L) and a high (Y_H) income group only. Let's further assume that everyone works and that 80 % earns a low income of 10,000 USD and the rest has high income earnings of 40,000 USD. As an example, we consider an inequality-averse society with a given concave social welfare function in income as in Eqn. (2), and resulting non-normalized welfare weights.¹¹

Table 1 Example of welfare weights under different normalizations

ρ	Inequality Aversion			
	0.0	0.5	1.0	1.5
	Normalization 1: $\omega_j^1 = \left(\frac{Y_j}{Y_L}\right)^{-\rho}$			
ω_L^1	1.00	1.00	1.00	1.00
ω_H^1	1.00	0.50	0.25	0.13
	Normalization 2: $\omega_j^2 = \left(\frac{Y_j}{\bar{Y}}\right)^{-\rho}$			
ω_L^2	1.00	1.26	1.60	2.02
ω_H^2	1.00	0.63	0.40	0.25
	Normalization 3: $\omega_j^3 = \theta(\rho)Y_j^{-\rho}$			
ω_L^3	1.00	1.11	1.18	1.21
ω_H^3	1.00	0.56	0.29	0.15

¹⁰ This solution has been proposed by Prof. B. Jacobs.

¹¹ For expository reasons the values for ρ have been stretched out a bit compared to the values stated in section 3.1. In chapter 4 we will revert to the values found in section 3.1.

Table 1 reports normalized welfare weights for three different normalizations and increasing societal inequality aversion, respectively $\rho = 0.0$, $\rho = 0.5$, $\rho = 1.0$ and $\rho = 1.5$. The first normalization uses the lowest income group (Y_L) as reference income: $Y_r = Y_L$. The second normalization uses the average income (\bar{Y}) as reference income: $Y_r = \bar{Y}$. Note that the average income (16,000 USD) has a welfare weight of one. In contrast, the third normalization is a weighted average normalization of the marginal weights to sum up to one, as earlier introduced by Eqn. (9). This normalization does not have a fixed reference income.

Only when society is not inequality averse, any normalized welfare weight is equal to one. The second column of Table 1 illustrates this. If welfare weights are normalized with respect to the lowest income group, the lowest income group is assigned a welfare weight of one, and any higher income is assigned a lower welfare weight. This is shown in the third row of Table 1: $\omega_H^1 < 1$ if $\rho > 0$. This normalization will only reproduce the results of a standard CBA if either society is not inequality averse, or if all benefits and costs of a public policy accrue to the lowest income group. For any other case, the benefits and costs of higher incomes will be valued less. This means, for example, that if a public policy has net benefits for any income class, the NPV of a distributionally weighted CBA with Normalization 1 will always be lower than in a standard CBA with equal weights of one.

The results of the second normalization illustrate that normalized welfare weights are larger than one for incomes lower than the average income ($\omega_L^2 > 1$ if $\rho > 0$; $Y_L < \bar{Y}$), and welfare weights are smaller than one for incomes higher than the average income ($\omega_H^2 < 1$ if $\rho > 0$; $Y_H > \bar{Y}$). This normalization, however, does not require that the weighted average of the applied welfare weights equals one. This implies that this normalization does not reproduce standard CBA outcomes when the net benefits are distributed equally across the population.

In contrast, the third normalization requires that standard CBA outcomes are reproduced when the net benefits are distributed equally across the population. To this end, constants $\theta(\rho)$ are computed for the considered degrees of inequality aversion. This implies that there is no fixed reference income.

Consider a policy that yields a total of 200 M USD for low incomes and 100 M USD for high incomes. Table 2 reports numerical results of distributionally weighted benefits for this case. The differences in outcomes are explained by the differences in the applied normalized welfare weights in Table 1.

Table 2 Corresponding distributionally weighted benefits (Mln USD)

Inequality Aversion: ρ	0.0	0.5	1.0	1.5
Normalization 1: $Y_r = Y_L$	300	250	225	213
Normalization 2: $Y_r = \bar{Y}$	300	316	360	430
Normalization 3: $\omega_j^3 = \theta(\rho)Y_j^{-\rho}$	300	278	265	258

Table 2 illustrates that distributionally weighted CBA outcomes may differ under different normalizations. Given the rate of inequality aversion, different normalizations yield wildly different outcomes. For example for an equality aversion rate of one ($\rho=1$), benefits may range from 225 to 360 depending on the normalization and this can be therefore higher or lower than the benefits according to the Hicks-Kaldor criterion ($\rho=0$).

If differentiated weighting is applied at all, normalizations 1 or 2 are commonly used. But they amount to a different scaling of the average dollar gain.¹² As a consequence, the results from different normalizations cannot be compared. They reflect rescaled objectives: a dollar to the poor is either considered to be worth a dollar under Normalization 1, or more than a dollar under the other normalizations. Which normalization is preferred may be interpreted as an issue of personal taste, as the ranking of projects for a given value of inequality aversion is independent of the normalization. However, the example shows that one has to be aware of the considered normalization for the interpretation of the results of a distributionally weighted CBA, as the numerical results in dollar terms may look very different for different normalizations.

From the above, we conclude that an application guideline for the normalization is needed to safeguard consistent application of marginal welfare weights and for comparability of results of different CBA's. In the case study application in Section 4, Normalization 3 is applied. This normalization has the property that if a project has equal net benefits for all income groups and therefore does not redistribute total benefits will always coincide with the unweighted Hicks-Kaldor case regardless of the degree of inequality aversion. As a consequence, only for this Normalization the results can be compared to the unweighted Hicks-Kaldor case and that any difference can be interpreted as only the result of a higher inequality aversion. Looking again at table 2 we can say that compared to the unweighted case the project's CBA outcome should be corrected downwards

¹² The reader may verify that these normalizations are transformations with $\omega_i^2 = \omega_i^1 c_1$, $\omega_i^3 = \omega_i^2 c_2$, and with $c_1 = \omega_L^2$ and $c_2 = \frac{1}{0.8\omega_L^2 + 0.2\omega_H^2}$.

by USD 22 mln if $p=0,5$, by USD 35 mln if $p=1$ and by USD 42 mln if $p=1,5$ for adverse inequality effects.¹³ These comparisons cannot be made for the other two normalizations.

Normalization 1 would also yield results that are at least easy to interpret. We do not advocate the use of average income as reference income, despite its use in literature applications and textbooks. The weighted average of the welfare weights does not add up to one, while it also lacks the compelling feature of Normalization 1 if there are only winners or only losers of a policy proposal.

3.3 Measuring redistribution effects

Sections 3.1 and 3.2 elaborated on marginal welfare weights (ω_i^n). We now turn to the welfare changes (ΔY_i) that result from a public policy across income classes. For some cases, a model including a breakdown by income groups can be used for this purpose. This could also be a combination of models, e.g. a micro simulation model with a micro data set on household income, a purchasing power model translating changes in gross income into net adjusted disposable income and a model for behavioral changes of labor supply. However, such a model or combination of models is not always available, at least not for most compilers of CBAs and not for many policy areas.

If such a model is not available, developing such a model and having access to the relevant microdata is complex, time-consuming and costly. For example, in the Netherlands, such a combination of models exists at the CPB Netherlands Bureau for Economic Policy Analysis. These models are specifically designed and used for forecasting and analyzing the impact of policy measures for social security and wage- and income tax. Substantial resources were allocated to develop these models. Other economic research institutes and private consultancy firms in the Netherlands do not have access to these models and do not have a set of similar models.

No such models may exist for other policy areas, for example for other taxes than wage and income taxes, such as VAT, corporation tax and excise duties, and for education policies and infrastructure. For other policy areas, assessing the net benefits for various income groups and attaching different weights to different income groups may also be less relevant for political decision-making. For example, for assessing the effects of a new road a breakdown by major types of stakeholders like those that will benefit by less travel time, those that pay for the road and those that are affected

¹³ The reader may check for himself that when dollar effects are equal for all (amounting to a total of USD 240 mln for the poor and a total of USD 60 mln for the rich), that normalization 3 will return a total CBA balance of USD 300 mln regardless of p . If the distribution of net benefits tilts further in favour of the poor, the project will see a positive adjustment for inequality effects compared to the unweighted Hicks-Kaldor balance.

most in terms of pollution and noise, is more meaningful and relevant than a breakdown of net benefits by income group.¹⁴

3.4 Accounting issues

Several accounting issues may need further investigation for real-world applications of marginal welfare weights. Some examples are:

1. Budgetary surplus redistribution and stakeholder accounting. For example, the revenues from a tax cut can be redistributed to consumers, but also to other stakeholders. Only consumer benefits are distributionally weighted. As a consequence, benefit accounting and surplus sharing is more important in a distributionally weighted analysis than in a standard CBA. This amounts to a more pressing need to account for incidence when distributional weighting is considered. This is complicated as it requires detailed general equilibrium accounting at the margin (not the average).
2. Marginal welfare weights can be estimated using the inverse optimal method, or from inferring from labour market reactions to policy initiatives. Both methods are limited to active labour market participants. Inactive groups such as retired or disabled people are not included. Within income brackets these groups may enjoy very different welfare weights compared to their working peers. This could affect the overall picture of income based social weighting differentials. It could also highlight a need to deal differently with people that can earn a living and those that are dependent.
3. Tax systems change over time. Marginal welfare weights from an inverse optimal model have to be updated after large tax reforms and over time.

4. Case study child care benefits

In this section, we analyze a complete and a partial repeal of child care benefits in the Netherlands. In the Netherlands, households with children in the age of 0-12 years are eligible for child care benefits if both parents or the single parent work. Net disposable household income¹⁵ is depicted in Figure 2. If child care benefits would be fully or partially repealed, this would have direct effects on the nominal adjusted disposable income of the beneficiary households.

¹⁴ Redistribution of monetary income and (re)distribution of other effects, like travel time and pollution, can in principle also be captured by introducing an extended income concept that includes non-monetary amenities and disamenities.

¹⁵ Adjusted for differences in size and composition of households.

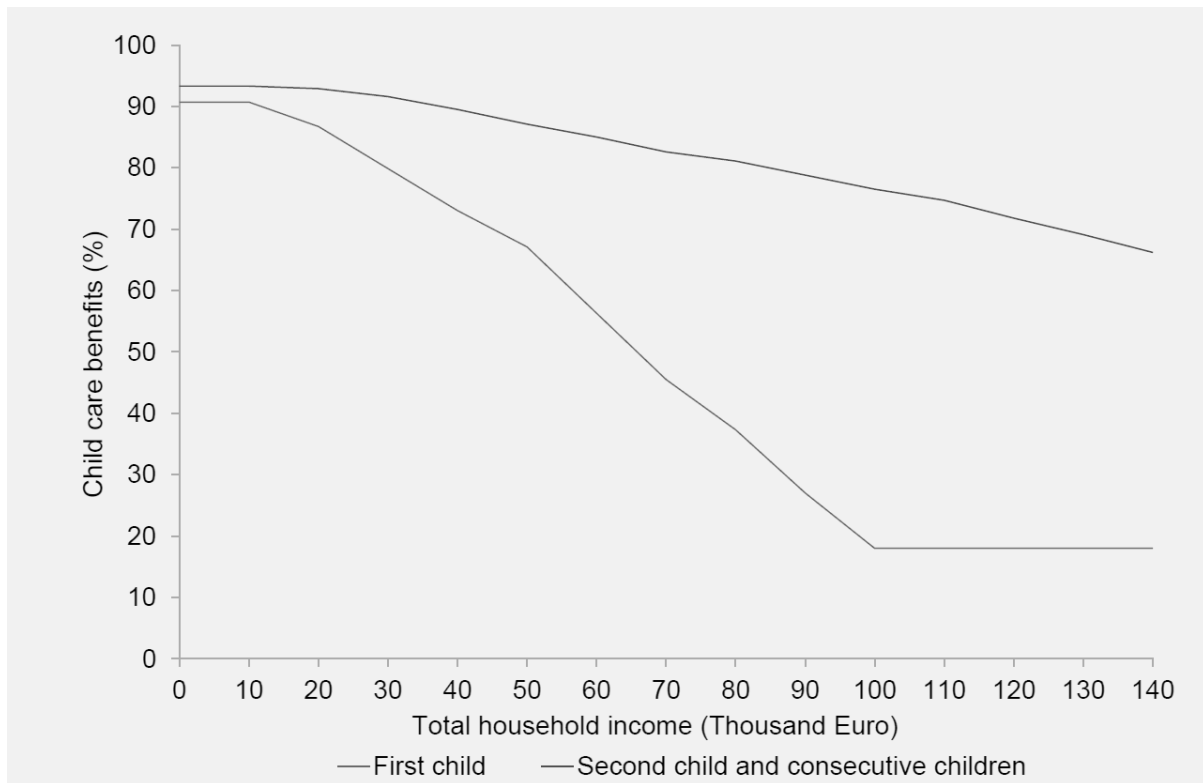


Figure 2 Percentage of private child care costs that are reimbursed as a function of total household income in 2014 (source: De Boer 2015, modified)

4.1 Complete repeal of child care benefits

The ex-ante effects of a complete repeal of child care benefits on household net adjusted disposable income are analyzed with microsimulation model MIMOSI (see Romijn et al. 2008 for a description). The results are reported in Table 3. For simplicity, nominal income changes have been aggregated in income classes of equal size, except for the lowest and highest income class. The lowest income class contains fewer observations. Some observations, for example students and certain other low incomes, have been deleted for practical reasons. This reduces reliability of the median estimates if the number of bins is increased. Alternatively, percentile estimates could be used.

The fourth column of Table 3 displays changes of disposable household income. Table 3 also shows that the child care benefits as share of household income is higher for lower incomes than for higher incomes. This reflects that for low income households the relative importance of child care benefits is generally higher than for high income households. The last column reports the totals (M Euro) of these income changes. For our analysis, these changes are weighted. Table 3 shows that a complete repeal would affect all beneficiary households negatively. Note that surplus redistribution is not studied for reasons explained earlier.

Table 3 Ex-ante effects of a complete repeal of child care benefits on disposable household income

N x 1000	Gross income categories (Thousand Euro)		Change in disposable income of households (%)	Total effect (M Euro)
	Minimum	Maximum	Median	
28.5	0.0	19.8	-24.3	-182
22.8	19.8	29.7	-19.3	-140
21.3	29.7	39.6	-15.7	-119
35.3	39.6	49.5	-10.4	-154
54.8	49.5	59.4	-9.5	-251
61.6	59.4	69.3	-7.7	-253
53.0	69.3	79.2	-6.9	-214
40.4	79.2	89.1	-5.7	-156
30.9	89.1	99.0	-4.9	-112
21.6	99.0	108.9	-4.0	-73
74.0	>108.9		-3.0	-246

Table 4 reports outcomes of the distributionally weighted income changes. We have normalized the welfare weights for the weighted average to sum up to one, i.e.: Normalization 3 as explained in Section 3.2. Note that the first row is equal the sum of the income changes in Table 3. The rest of the results in table 4 states that inequality aversion imposes an extra cost in terms of the CBA balance. In other words the measure increases only when there is inequality. And the effect becomes worse as society's inequality aversion becomes more pronounced. Apparently, discontinuing child care benefit as a whole is not a good idea and it becomes worse if you take into account inequality. Table 3 shows why: more than three quarters of the total child care benefits are received by households with gross incomes of 39600 Euro or more. These are given less weight in the analysis as compared to the benchmark of no inequality version, so their losses are discounted. But the losses of the low incomes are supercharged (see figure 1) resulting in a net inequality loss.

Table 4 Weighted totals of household income changes for a complete repeal of child care benefits

Inequality Aversion (ρ)*	Total Weighted Income Changes (B Euro)
0.00	-1.9
0.27	-2.0
0.54	-2.1
1.08	-2.3

* These values were established in section 3.1

The results in Table 4 depend on the normalization choice. For example, if we would normalize the welfare weights with respect to the lowest household income, i.e. the lowest income is assigned a weight of one and any higher income is assigned a marginal weight of less than one, then correction

for inequality for this case study would always be positive if society is inequality averse ($\rho > 0$) as compared to the benchmark ($\rho = 0$). Given that all household incomes are negatively affected and are assigned a welfare weight of one or less than one, the correction as compared to the benchmark has to be positive for this case and normalization 1. Thus, one has to be careful with the interpretation of the numerical results of a distributionally weighted analysis in dollars. The interpretation of these dollars depends on the normalization choice. However, the normalization choice does not affect the ranking of projects.

4.2 Partial repeal of child care benefits

A partial repeal of child care benefits for high-income households would yield better ex-ante results under inequality aversion than a complete repeal. Consider a proposal to repeal child care benefits for incomes of 39600 Euro or more. Child care benefits are kept unaltered for lower incomes.

Table 5 shows results for this partial repeal. Identical normalized marginal welfare weights are applied as in the previous analysis. Note that the total increases (becomes less negative) in inequality aversion. In other words, the correction for inequality is positive as compared to the benchmark of undifferentiated welfare weights (Hicks-Kaldor; $\rho=0$).

Table 5 Weighted totals of household income changes for a partial repeal of child care benefits

Inequality Aversion (ρ)	Total Weighted Income Changes (B Euro)
0.00	-1.5
0.27	-1.4
0.54	-1.2
1.08	-0.9

5. Conclusions and discussion

The conventional CBA practice of adding up benefits and costs irrespective of who gains or loses is based on the Hick-Kaldor compensation principle. This practice does not show the impacts of a policy proposal on the income distribution in the net benefits of a CBA. If society is averse to inequality this approach fails to account for this.

This paper elaborates on distributionally weighted CBA for real-world applications. We highlight its building blocks, including a marginal welfare weight selection, a weight normalization and benefit accounting choices, and an analysis of redistribution effects. The latter may be obtained from a purchasing power model for analysis of tax policy changes. The case study in this paper on a repeal

of child care benefits provides an example of the additional insights that are obtained with a distributionally weighted CBA. The results also illustrate why effects of inequality aversion may be easily misunderstood. We conclude that guidelines for marginal welfare weights are needed to safeguard consistent applications and comparability of CBA-case study results.

We emphasize that distributionally weighted CBA may be demanding, prone to misuse and might distract attention from the decision problem at hand. Purchasing power models may be available for tax policy analysis, but models to analyze detailed income redistribution effects may be absent for applications in other policy domains. For such cases, a stakeholder analysis may be performed, but distributionally weighted analysis of net welfare effects may be cumbersome. Moreover, marginal welfare weight selections might be misused to reverse-engineer results if application guidelines are lacking.

The selection of marginal social welfare weights remains problematic due to its normative nature.¹⁶ This paper applies smoothened marginal welfare weights from an inverse optimal taxation model. This has the advantage that anomalies in the semi-empirical results from such a model are removed. However, the presence of anomalies also reveals the absence of tax optimality. An inequality deflator, in contrast, is more general and does not require a social welfare function (Hendren 2014). Three different marginal welfare weights are derived from the inverse optimal taxation model for the Netherlands: the smoothened marginal welfare weight itself (0,54), a lower bound of half this weight (0,27) and an upper bound of double this weight (1,08).

The case study on child care benefits illustrates that a distributionally weighted CBA may provide complementary information to an analysis of income redistribution effects, for example by means of a purchasing power analysis. A distributionally weighted analysis aggregates income changes. The results of a purchasing power analysis, in contrast, contain more detailed effects of a policy on net adjusted income changes, for example with respect to household compositions. A purchasing power analysis usually reports relative income changes as percentages as well, which was shown for the case study on child care benefits. This is in fact a weighting exercise. For example, a double income halves the relative income change. However, these weights are implicit and relative income changes are reported but not weighted in such an analysis.

¹⁶ Standard CBA with unit welfare weights is equally normative, but has the advantage of simplicity and will only yield substantially different results when the distributional effects of the project or policy proposal are sufficiently large.

For further research, more applications are needed to study the topics addressed in this paper for different policy domains. For CBA practice, we advise to report net benefits or disposable income changes of public policies if this information supports decision-making and outweighs its transaction costs. Purchasing power analysis and stakeholder analysis provide tools to show distributive effects in CBAs. If these results are insufficient for informed decision-making, distributionally weighted CBA could be considered as a complementary tool for policies that involve major efficiency-equity tradeoffs.

References

- Boardman, A. E., Greenberg, D. H., Vining, A. R., and Weimer, D. L. (2006). *Cost-Benefit Analysis: Concepts and practice. 3rd Edition*. Upper Saddle River: Pearson.
- Bos, F., van der Pol, T.D. and Romijn, G. (2017). Should CBA's include a correction for the marginal excess burden of taxation?, paper presented at the 9th Conference of the Society for Benefit-Cost Analysis, March 15-17, 2017, Washington DC. Forthcoming as CPB Discussion Paper.
- Cowell, F.A. and K. Gardiner (1999). Welfare Weights. STICERD, London School of Economics, Economics Researchpaper 20.
- Dahlby, B. (2008). *The marginal cost of public funds: Theory and applications*. Cambridge: MIT Press.
- De Boer, H. W. (2015). *Onderzoek alternatieve vormgeving kinderopvangtoeslag (in Dutch)*. Den Haag.
- Evans, D. J. (2005). The elasticity of marginal utility of consumption: Estimates for 20 OECD countries. *Fiscal Studies*, 26(2), 197-224
- Florio, M. (2014). *Applied welfare economics: Cost-Benefit Analysis of projects and policies*. New York: Routledge.
- Harberger, A. C. (1978). On the use of distributional weights in Social Cost-Benefit Analysis. *Journal of Political Economy*, 86(2), S87-S120
- Hendren, N. (2014) The Inequality Deflator: Interpersonal comparisons without a social welfare function, NBER Working Paper No. 20351
- HM Treasury (2011). *The Green Book: Appraisal and evaluation in central government treasury guidance*. London.
- Jacobs, B. (2016). The marginal cost of public funds is one at the optimal tax system, mimeo: Erasmus Universiteit Rotterdam.
- Jacobs, B., E.L.W. Jongen and Floris T. Zoutman (2017). Revealed Social Preferences of Dutch Political Parties, *Journal of Public Economics*, forthcoming.
- Kind, J., Botzen, W. J., and Aerts, J. C. J. H. (2016). Accounting for risk aversion, income distribution and social welfare in cost-benefit analysis for flood risk management. *Wiley Interdisciplinary Reviews: Climate Change*. doi: 10.1002/wcc.446
- Layard, R., S. Nickell and G. Mayraz (2008). The marginal utility of income. *Journal of Public Economics*, 92, 1846-1857.
- Koopmans, C., Heyma, A., Hof, B., Imandt, M., Kok, L., and Pomp, M. (2016). *Werkwijzer voor kosten-batenanalyse in het sociale domein: Hoofdrapport*. Amsterdam.
- Lockwood, B. B., and Weinzierl, M. (2016). Positive and normative judgments implicit in U.S. tax policy, and the costs of unequal growth and recessions. *Journal of Monetary Economics*, 77, 30-47. doi: <http://dx.doi.org/10.1016/j.jmoneco.2015.10.006>
- Pearce, D. and D. Ulph (1995). A Social Discount Rate for the United Kingdom, CSERGE Working Paper No 95-01, School of Environmental Studies University of East Anglia Norwich.

- Romijn, G., Goes, J., Dekker, P., Gielen, M., and van Es, F. (2008). *MIMOSI: Microsimulatiemodel voor belastingen, sociale zekerheid, loonkosten en koopkracht (in Dutch)*. The Hague.
- Romijn, G., and Renes, G. (2013). *General Guidance for Cost-Benefit Analysis*. The Hague.
- Werkgroep Kosten van belastingheffing en MKBA's. (2016). *Kosten van belastingheffing, baten van inkomensherverdeling en MKBA's (in Dutch)*. The Hague.
- Zoutman, F. T., Jacobs, B., and Jongen, E. L. W. (2016). Redistributive politics and the tyranny of the middle class. T. Institute.



Publisher:

CPB Netherlands Bureau for Economic Policy Analysis
P.O. Box 80510 | 2508 GM The Hague
T (088) 984 60 00

October 2017