

Which Households Matter Most?

Generalised Social Marginal Welfare Weights and Indirect Tax Reform

David Madden¹ and Michael Savage^{1,2}

¹*University College Dublin, Ireland*

²*Economic and Social Research Institute, Dublin, Ireland*

April 14, 2014

Abstract

Social marginal welfare weights play an important role in areas of applied public policy analysis such as tax reform. These weights reflect the values of the social planner, or equivalently the underlying social welfare function. Thus when evaluating the welfare effects of a tax reform, different weights may be assigned to different groups in society, with the weights reflecting the relative importance of these groups in terms of aggregating gains and losses arising from the tax reform. Until recently, the majority of research has adopted social marginal welfare weights based upon a Utilitarian approach, with concave utility of income functions, so that the poorest agent (individual or household) in society receives the highest welfare weight and weights decline as income increases. However, a number of recent papers have questioned the Utilitarian-based approach and have suggested potential alternatives (Saez and Stantcheva (2013), Weinzierl (2013)). Many of these papers exploit specifically designed online surveys which elicit respondents' preferences for redistribution (including the two papers cited here). However, there are few examples applying these alternative weighting schemes to traditional, nationally representative, datasets, and in particular, few comparisons of how these alternative weighting schemes would affect the distribution of the welfare effects of a specific tax reform in comparison to the Utilitarian-based approach.

This paper aims to fill that gap. Using the nationally representative 2009/10 Irish Household Budget Survey, we apply a range of alternatives to Utilitarianism in determining the distribution of social marginal welfare weights, and compare these distributions to that arising from the traditional Utilitarian approach. The alternative weighting schemes we analyse are based upon: the principles of Equal Sacrifice, poverty alleviation, government self-interest and finally the redistribution of "luck" income. The distribution of welfare weights arising from these approaches are found to differ appreciably from the distribution based upon Utilitarian weights. A simple indirect tax reform model is then estimated and applied to the different distributions of welfare weights to investigate the sensitivity of tax reform recommendations to these distributions. Given the importance of social marginal welfare weights in areas of public policy analysis such as optimal labour and commodity tax design, and tax

reform evaluation, we believe this detailed examination of the alternatives to Utilitarianism, and their application to a household budget survey dataset, is an important addition to the literature.

1 Introduction

In evaluating public policies such as optimal taxes and transfers, the standard approach is what has become known as the *welfarist* approach. It is assumed that the government wishes to maximise a social welfare function, and that this social welfare function depends upon individual utilities, which are in turn represented by utility functions. Social welfare is maximised subject to a budget constraint and account is also taken of behavioural responses to policies.

Seminal papers using this approach include Diamond and Mirrlees (1971) in the area of commodity taxation and Mirrlees (1971) in the area of income taxation. This approach is also consistent with the literature on marginal tax reform (Ahmad and Stern 1984) where, instead of trying to locate the optimum, instead we search for marginal welfare-improving, tax-neutral reforms (though the social welfare function is typically not specified in this approach, since rather than trying to maximise social welfare, we are trying to identify local, welfare-improving reforms). A small change in the tax on good i , dt_i , will lead to a welfare change dW and also a change in revenue dR . At the optimum the ratio of dW/dR for each small tax change should be the same. If not, then a welfare-improving marginal tax reform is possible.

In checking for the existence of such tax reforms, analysts must calculate dW , the change in social welfare following the change in the tax on good i . As a first-order approximation, the change in welfare is given by the sum of each household's actual consumption of good i (following Roy's identity) multiplied by its marginal social welfare weight. It is the derivation of these welfare weights which is the principal focus of this paper.

Using the welfarist approach referred to above, welfare weights are typically declining (or at least not increasing) in whatever measure of household resources is employed in the analysis (e.g. income or expenditure). Welfare weights which are declining in income are justified via the declining private marginal utility of income.

In a recent contribution Saez and Stantcheva (2013) point out that the welfarist approach is merely one amongst a number of approaches which can be used to generate marginal social welfare weights. They propose instead generalised marginal social welfare weights. These weights simply reflect the relative value of marginal consumption which society places upon each individual and can be used to evaluate marginal tax reforms. The implicit optimum with these welfare weights does not necessarily maximise an ex ante social welfare function. Instead, the system is optimal if no small revenue neutral reform leads to a net increase in social welfare, where gains and losses in social welfare arising from the tax change are weighted by the generalised social welfare weights.

Different approaches will generate different sets of generalised social welfare weights and candidates for these different approaches would include perspectives such as Libertarianism, Equality of Opportunity etc. As their name suggests, such generalised welfare weights nest as a special case the welfarist approach, while also accommodating other approaches. Accommodating perspectives other than welfarism is important, as empirical evidence suggests that views on taxation and redistribution typically take into account factors other than those implicit in welfarism. Such factors can include whether the process generating the income

distribution is perceived as being fair and also more detailed information on the identity of a person apart from their income (e.g. if their net tax payment is positive). The latter factor is assumed away under the “anonymity” principle usually adopted in welfarism. Welfarism also takes no account of the source of income and makes no distinction between “fair” and “unfair” income.

Generalised welfare weights share with the marginal tax reform analysis of Ahmad and Stern the feature that the social welfare function is not explicitly specified. Instead, local tax reforms are evaluated and the welfare effect will depend upon the welfare weights chosen. These welfare weights could be generated from individual utility functions and a social welfare function (since welfarism is nested with the generalised approach) but they could also be generated using other principles which more directly reflect society’s views on justice, such as those listed above. As explained by Saez and Stantcheva, the analysis of such reforms is likely to be much closer to actual decisions taken by policy-makers, as opposed to the design of an optimal system. As pointed out by Ahmad and Stern (1984), it is also typically less informationally demanding.

Thus marginal tax reform analysis is ideal for examining the sensitivity of recommended reforms to the choice of how welfare weights are generated, and in particular to the implications of departing from welfarism. That is the contribution of this paper: we carry out marginal tax reform analysis on the Irish indirect tax system, using generalised welfare weights including both the welfarist and non-welfarist approaches. The non-welfarist approaches we adopt are to some extent limited by data availability, but we are still able to incorporate a reasonably wide range of generalised welfare weights. We also confine ourselves to cases with non-negative welfare weights, otherwise we would violate the constrained Pareto efficiency property. To the best of our knowledge this is one of the first attempts to empirically examine the sensitivity of actual tax reforms to a wider range of welfare weights than those implied by welfarism.

It is perhaps worth stressing at this stage that we are not advocating the use of any particular set of generalised welfare weights, whether arising from welfarism or any other approach. Our goal is merely to empirically investigate how sensitive marginal tax reform recommendations are to the choice of welfare weights, when the menu of welfare weights extends beyond welfarism (for examples of how marginal reforms vary with respect to welfare weights within the welfarist approach, see Ahmad and Stern (1984) and Madden (1995)).

2 Related Literature

When evaluating the welfare effects of a tax reform, different weights may be assigned to different groups in society, with the weights reflecting the relative importance of these groups in terms of aggregating gains and losses arising from the tax reform. Until recently, the majority of research has adopted social marginal welfare weights based upon a Utilitarian approach, with concave utility of income functions, so that the poorest agent (individual or household) in society receives the highest welfare weight, and weights decline as income increases. A parameter indicating the elasticity of the social marginal utility of income is often used to determine the rate of decline in the weights¹ (see, for example, Saez (2001), Decoster and Schokkaert (1990), Madden (1995), Ray (1986)). Weinzierl (2012, 2013) questioned this “default” assumption of Utilitarianism. In a US based study, he presented evidence of strong

¹see Section 4.2 for more detail

support for tax systems based at least in part on the principle of Equal Sacrifice rather than on pure Utilitarianism. Using a specifically designed online survey, Weinzierl (2013) asked respondents' opinions on a number of hypothetical tax and benefit systems. The hypothetical systems ranged from one based on Equal Sacrifice, whereby each household paid 14 per cent of their income in tax, to one based on Utilitarianism, whereby the average tax rate ranged from -731 per cent for the poorest household to 79 per cent for the richest household. In a straight choice between these two extremes, he found that nearly 60 per cent of respondents preferred the Equal Sacrifice system to the Utilitarian system. Offering tax and benefit systems with a combination of the two principles, he found an even larger majority preferred a system with at least some Equal Sacrifice characteristics.

Variations exist in the interpretation of the principle of Equal Sacrifice. The Equal Sacrifice tax system used in Weinzierl (2013) imposes the same average tax rate on all individuals regardless of income, which follows the principle of Equal *Proportional* Sacrifice. As Young (1987) suggests, Equal proportional sacrifice means that everyone foregoes the same percentage of utility in paying taxes. An alternative interpretation is Equal *Absolute* Sacrifice, whereby everyone foregoes the same amount of utility in paying taxes. Young shows that even with certain concave utility functions, the principle of Equal Sacrifice, in particular Equal Absolute Sacrifice, can result in regressive tax systems.

While Weinzierl did not directly translate his results into a profile of social marginal welfare weights, Saez and Stantcheva (2013) used a similar specifically designed online survey to elicit social preferences, also from US respondents. Their approach was to replace the "standard" Utilitarian welfare weights with generalised social marginal welfare weights which "directly reflect society's view for justice". The generalised approach departs from the Utilitarian framework as the welfare weights are no longer based on underlying individual utility functions or social welfare functions. They argued that the generalised welfare weights solve many of the puzzles of the Utilitarian approach (detailed in Section 4.2.1), while it also has the ability to capture a number of prominent alternatives to Utilitarianism, such as poverty alleviation, Libertarianism and redistribution based on the nature of the income in a household, all of which are ignored in the Utilitarian approach. Also using a specifically designed online survey, Kuziemko et al. (2013) split respondents into treatment and control groups to examine the effect of an informational treatment on preferences for redistribution. In the study, the treatment included providing information on the distribution of income in the US, each respondent's position within that distribution, and correlations between top tax rates and economic growth in the US. When comparing preferences for redistribution across social class, they found that support for transfer programs for the poor declined among low income and low education groups after the informational treatment. Thus, even when low income or education respondents were fully informed on the distribution of income in the US, support among this group for Utilitarianism (which would encourage transfers to the less well-off) was low. The ability to capture alternative social preferences becomes more important in this respect.

Of course, it is not only since 2013 that the limitations of the Utilitarian approach have been recognised. Mirrlees (1974) highlighted shortcomings of the Utilitarian approach, and allowed for agents with different skill levels to "envy" each other based on the number of hours they have to work. Elsewhere, Fleurbaey and Maniquet (2007) allowed for "fairness" in optimal income taxation, by defining social preferences that differentiate between redistribution based on earning ability and preferences for work and leisure. Their model supports redistribu-

tion based on inequality between agents who differ in their earning ability, but have identical preferences. However, redistribution between agents with the same earning ability, but differing taste for work, is discouraged. Remaining within the Utilitarian framework, Alesina and Angeletos (2005) allowed individual utility to decline in the amount of “unfair” income in society by including an extra term in the individual utility functions. Saez and Stantcheva (2013) cautioned against this approach however due to the potential for non-standard individual behaviours, as well as the fact that social preferences may not be respectful of individual preferences.

A move outside of the Utilitarian framework does not automatically result in a profile of welfare weights substantially different to that arising from the traditional Utilitarian approach. In an extensive review of the optimal taxation literature, Piketty and Saez (2012) argued that that relative income concerns are a much more powerful and realistic way to justify social welfare weights decreasing with income than standard Utilitarianism with concave utility of income. They argued that specifying the social welfare weights in such a manner captures the idea that a decrease in one person’s income may increase the utility of others. The relatively low welfare weight on a richer person would therefore capture the “external effect” of a decrease in his income increasing the utility of others.

In determining the profile of social marginal welfare weights in tax reform models, the literature reviewed here shows that a number of alternative approaches to Utilitarianism are possible. While several authors have suggested extensions or alternatives to Utilitarian welfare weights, many of the empirical analyses exploit specifically designed online surveys which elicit respondents’ preferences for redistribution. There are few examples applying these alternative weighting schemes to traditional, nationally representative, datasets, and in particular, few comparisons of how these alternative weighting schemes would affect the distribution of the welfare effects of a specific tax reform in comparison to the Utilitarian-based approach.

We believe this paper makes two contributions to the tax reform literature. First, we examine how the distribution of welfare weights changes as we move away from Utilitarianism. While much of the recent literature has used online surveys to explore what alternatives to Utilitarianism may exist, the use of a traditional, nationally representative dataset in this paper allows us to identify changes in the relative weight given to different agents across the income distribution. Second, using these profiles of welfare weights, we examine the extent to which a specific public policy (an indirect tax reform) is sensitive to the alternative to Utilitarianism. Given the importance of social marginal welfare weights in areas of public policy analysis such as optimal labour and commodity tax design, and tax reform evaluation, we believe this detailed examination of the alternatives to Utilitarianism, and their application to a specific policy proposal, is an important addition to the literature.

3 Data

The primary source of data used in this paper comes from the 2009/10 wave of the Irish Household Budget Survey (HBS). The survey includes detailed information on household expenditure, income and tax payments, as well as socio-economic variables on 5,891 households. For the purposes of the analysis we trim the bottom 1 per cent of observations according to household gross income. This eliminates 14 cases of 0 gross income, and also a further 39 cases with household gross income (weekly, unequivalised) of less than €100. A number of the sets

of welfare weights in the following sections are dependent on income and expenditure values in the extremes of the distributions. Removing the bottom 1 per cent reduces the possibility of the marginal welfare weights being significantly affected by outliers in the data. Specific details of the components of the HBS data are given in the relevant sections.

3.1 Income Data

Commonly the first step to determining the social marginal welfare weight of a particular household, in the Utilitarian case in particular, is to measure the income of that household. While the primary purpose of the HBS is to measure household expenditure, the data also reports detailed information on income. For each household, direct, gross and disposable income levels are reported, as well as more detailed information on income source, such as level of employment income, state transfer income and income from investments and properties. Following the most common approach taken in distributional analysis literature, we choose disposable income as our main measure of household resource.

Given the HBS is conducted at household level, all income figures are reported for the household as a single unit. Individual income levels therefore cannot be identified. This raises two points regarding the analysis in the following sections that are worth noting. Firstly, the decision to equivalise income can have significant implications for distributional analysis. According to Creedy (2006), the use of equivalence scales involve difficult value judgements for economists. Indeed, Coulter et al. (1992) showed that not only can the decision to equivalise significantly alter relative incomes across household types, but the choice of the scale itself also affects relative incomes. In relation to inequality and poverty measures, they found that changes induced by the choice of equivalence scale are “not small”. Examining in detail the choice of between equivalence scales is beyond the scope of this paper. We follow the approach used by the Central Statistics Office (CSO) of Ireland when reporting income distribution and poverty statistics. They, and we, use a scale of 1 for the first adult, 0.66 for subsequent adults, and 0.33 for children to equivalise income. This is also the scale implicit in many of the main welfare payments in Ireland, where adult dependant allowances are approximately two-thirds of the main payment, and child dependent allowances are approximately one-third of the main payment. Secondly, using household income to measure the resource of individuals within that household implies full income sharing among household members. While evidence exists that full income sharing in a household may not hold in certain circumstances (see Watson et al. (2013), Lundberg et al. (1997)), the nature of the data means we are unable to identify individual level income within a household. As Cowell and Gardiner (1999) suggest, it is common in the literature to make the pooling, or equal-sharing, assumption.

3.2 Direct and Indirect Tax Data

As well as income data, we also require information on tax payments of each household in some of the alternative welfare weight specifications. The HBS identifies household income in three key stages: direct income, gross income and disposable income. Gross income is simply direct income plus state transfer payments. Disposable income is gross income minus income tax and social insurance contributions. It is straightforward therefore to identify each household’s direct tax and social insurance contributions, as well as any state transfer payments received by anyone in the household.

The indirect tax payments made by a household are less straightforward to identify². In order to do so, we assign a VAT rate to each of the expenditure components in the HBS³. Each item is assigned either the standard rate (21 per cent⁴), the reduced rate (13.5 per cent) or a 0 per cent rate (items can have a 0 per cent VAT rate applied or be exempt from VAT). For goods where excise duties apply (alcohol, tobacco and fuel), we rely on statistics from the Revenue Commissioners (2012) which report the tax component of the price of a sample good from each of these categories. For example, Revenue report the tax component of various quantities of lager, stout, cigarettes, petrol and diesel, among others, which we apply to the relevant expenditure categories in the HBS⁵. From the estimates of the VAT and excise duties paid by each household, we can estimate each household's indirect tax contribution. Table 1 shows the direct tax payments, state transfer payments, and indirect tax payments as a percentage of gross income in each decile⁶. The bottom deciles receive the largest share of their gross income through state transfer payments, with state transfers making up over 80 per cent of gross income in the second decile. The direct tax system (including social insurance contributions) is highly progressive, with the average tax rate increasing in each decile. A regressive pattern emerges from the indirect tax payments, with indirect taxes accounting for a quarter of gross income in the poorest decile, but only 5.5 per cent in the top decile⁷.

Table 1: Tax and Welfare Payments as Percentage of Gross Income by decile

Decile	State Transfers	Direct Tax and SSCs	Indirect Taxes
<i>As % of Gross Income</i>			
Bottom	70.1	1.2	25.0
2	81.5	1.6	16.7
3	73.3	2.1	14.9
4	58.3	3.1	13.4
5	43.2	5.3	12.7
6	28.5	7.9	11.6
7	15.3	11.7	9.9
8	10.6	13.7	9.2
9	6.0	17.7	8.2
Top	2.2	22.9	5.5

²We measure only the VAT and excise duty payments made by a household. Additional indirect tax payments, such as the air travel tax, are omitted. See Collins and Turnbull (2013) for a more detailed method of estimating indirect tax payments.

³We are grateful to Sen Lyons for help identifying the relevant VAT rate on each expenditure component in the HBS.

⁴The standard VAT rate was reduced from 21.5 per cent to 21 per cent in Budget 2010. We assume the 21 per cent rate applied to all expenditure in the survey, although some will have been taxed at the higher 21.5 per cent.

⁵For example, in the HBS expenditure is reported for spirits consumed in the home. The Revenue Commissioners report the tax component of the price of a bottle of whiskey. We therefore apply the tax rate on a bottle of whiskey to all spirits consumed in the home.

⁶The two tax columns can be interpreted as average tax rates (direct and indirect) for each decile.

⁷Collins and Turnbull (2013) find a very similar pattern of tax payments by decile. There are slight discrepancies between the figures here and in their study due to small differences in methodology, and the inclusion of indirect taxes other than VAT and excise in Collins (2013). Overall the patterns found in both studies are very similar.

4 Social Marginal Welfare Weights

In this section, we use the HBS data discussed above to examine the distribution of social marginal welfare weights based on a number of alternative criteria. We follow Creedy's (2006) advice that "the main contribution economists can make is to examine the implications of adopting a range of alternative value judgements". Rather than attempting to identify the profile of weights that policy makers in Ireland should or do use in tax and benefit design, our aim is simple - to make clear the implications for tax reform analysis of using different criteria to determine who is most "deserving" in society.

4.1 Concentration Curves

We present the distribution of each set of welfare weights, $w(\cdot)$, by concentration curves according to equivalised disposable income⁸. This approach allows us to highlight a number of features of each distribution in one graph. We present the concentration curves relative to the 45 degree line, which represents the case where the welfare of each household is valued equally. Each point on the curve indicates the percentage share of social marginal welfare weights (y-axis) assigned to the corresponding proportion of the population, ranked by equivalised disposable income (x-axis)⁹. In each case, the welfare weights are normalised so that the household with the highest weight has a weight of one.

Concavity: A strictly concave¹⁰ distribution of welfare weights indicates that lower income households have relatively high marginal welfare weights. For example, the concentration curve for the Utilitarian weights (with any positive level of inequality aversion) based on equivalised disposable income is, by design, concave in equivalised disposable income. A strictly convex concentration curve¹¹ indicates that low income households have relatively low marginal welfare weights.

Slope: Related to the degree of concavity is the slope at each point of the concentration curve. The slope of the concentration curve gives the rate of change of the cumulative share of weights with respect to the cumulative share of income. In other words, the slope of the ray tells us the ratio of cumulative share of weights to cumulative share of income. Therefore, a section of the concentration curve with relatively steep slope (more than 45 degrees) indicates that that portion of households have a larger than proportional share of marginal welfare weights. A more shallow slope (less than 45 degrees) indicates a low concentration of welfare weights for that portion of households.

Area between Concentration Curve and 45 Degree Line: The larger the area between the concentration curve and the 45 degree line, the more unequal the distribution of welfare weights across the income distribution. As before, if the concentration curve lies above the 45 degree line, then the distribution of welfare weights is concentrated among poor households. The further the curve is above the 45 degree line, the more concentrated are the weights among the poor. If the concentration curve lies below the 45 degree line, the inequality is caused by a larger than proportional share of weights among the higher income households. Again,

⁸We follow O'Donnell et al.'s (2008) method of graphing concentration curves in Stata.

⁹Having disposable income on the x-axis can be seen as setting the Utilitarian case with disposable income as the measure of resource as our 'base' case. The concentration curves then show how each alternative differs from this base case.

¹⁰Concave in income. Concentration curve above the 45 degree line; increasing at a decreasing rate

¹¹Convex in income. Concentration curve below the 45 degree line; increasing at an increasing rate

the larger the gap between the curve and the 45 degree line, the more concentrated are the weights among the rich.

Dominance: The concentration curve approach makes it straightforward to compare the degree of equality across the income distribution in two or more distributions of welfare weights, particularly when all curves of interest lie on the same side of the 45 degree line. If the concentration curve for one set of welfare weights lies closer to the 45 degree line than another at all points of the distribution, the first curve is said to dominate the second and the ranking by degree of inequality is unambiguous. Where concentration curves cross, we cannot unambiguously determine which distribution of weights is more equal across the income distribution than another simply by examining the graph. This approach is similar to Lorenz curve analysis. However, unlike Lorenz curves, the welfare weight concentration curves can lie above or below the 45 degree line¹². This results in difficulties in comparing curves that lie on opposite sides of the 45 degree line, or when a curve crosses the 45 degree line. For this reason, we include a table of concentration indices in an appendix¹³.

4.2 Utilitarian Social Welfare Weights

We begin by outlining the specification of social marginal welfare weights under the standard Utilitarian approach. In this case, we have a social welfare function $G(u^h)$ which is a function of household utilities alone. Social welfare is simply a weighted sum of household utilities. The key feature of the Utilitarian approach is that for any positive value of inequality aversion, the weight placed upon the welfare of household h is declining in some measure of household resource.

Formally we can specify the welfare weights as per a commonly used utility of income function due to Atkinson (1970), where household utility, U^h , is a function of household resource alone:

$$U^h(I^h) = \frac{k(I^h)^{1-e}}{1-e} \quad \text{if } e \geq 0 \text{ and } e \neq 1 \quad (1)$$

$$U^h(I^h) = k \log(I^h) \quad \text{if } e = 1 \quad (2)$$

where I^h is the measure of resource of household h , and k is chosen for normalisation. e is a parameter reflecting inequality aversion, and is the key parameter in determining relative welfare weights in the Utilitarian setting. For any $e > 0$ we have welfare weights satisfying strict concavity, so that as household income increases, the welfare weight of that household strictly decreases; the larger the parameter e , the faster the weights decline in income. As we assume a constant inequality aversion parameter¹⁴, for every percentage point increase in a household's income, the welfare weight of that household decreases by e per cent.

In practice, we normalise the welfare weights so that the poorest household has $\beta^h = 1$. We therefore measure the welfare weight of household h relative to the poorest household:

$$\beta^h = \left(\frac{I^1}{I^h} \right)^e \quad (3)$$

¹²Lorenz curves graph income share from poorest to richest agent, so that the concentration curve must lie below the 45 degree line

¹³We calculate the concentration index using O'Donnell et al.'s (2008) "convenient covariance" method.


¹⁴Which is equivalent to assuming a constant elasticity of marginal utility of income

where I^1 is the resource of the poorest household. Higher values of e result in higher relative weight on the welfare of the poorest household.

Much like Donaldson and Weymark's (1980) *S-Gini* measure of inequality, Atkinson's utility functions allow a range of social preferences to be captured by varying one parameter. The Benthamite, or extreme Utilitarian, case exists where $e = 0$, so society has no aversion to inequality, and each household has $\beta^h = 1$. Social welfare is an therefore unweighted sum of household utilities. As e moves to infinity, we get closer to the Rawlsian case, where only the welfare of the poorest household is taken into account. The Benthamite and Rawlsian cases can be seen as the two limiting cases of the Utilitarian approach. Changing the value of e will alter the relative welfare weights in a given distribution, but will not alter the ranking of households.

The "correct" value of e is an empirical matter, beyond the scope of this paper¹⁵. Figure 1 compares the distribution of marginal welfare weights across the income distribution, with I^h equal to equivalised disposable income (giving set of social marginal welfare weights $w(y_d)$), as we allow e to range between zero and five¹⁶. As expected, as e becomes larger, the weights become more and more concentrated in the lower end of the income distribution. With $e = 5$, we are close to the Rawlsian case where only the welfare of the poorest household matters.

Figure 1: Concentration Curve for Utilitarian Welfare Weights



4.2.1 Beyond Utilitarianism

While Utilitarian welfare weights can be used to capture a range of social preferences, a number of authors have suggested moving beyond the Utilitarian framework when determining the relative weight placed on the welfare of a particular household. Saez and Stantcheva (2013)

¹⁵In their guide on Appraisal and Evaluation in the UK Central Government, H.M. Treasury (2011) assume a value of $e = 1$ and measure utility as $U = \log C$. As far back as 1939, former UK Chancellor High Dalton suggested that e lay between 1 and 2 (Dalton 1939).

¹⁶The 45 degree line can be interpreted as the case where $e = 0$.

identify a number of limitations of the Utilitarian approach. First, if individuals do not respond to taxes, Utilitarianism recommends a 100 per cent tax and full redistribution¹⁷. Second, Utilitarianism does not discriminate on the type of income, so that redistribution of “deserved” and “luck” income is equally desirable. Similarly, no distinction is made between “free loaders” (those who would work in the absence of benefits) and “deserving poor” (those who are unable to work) in the Utilitarian framework. Finally, the use of tags can be optimal¹⁸. More generally, within the Utilitarian framework, only information that enters the household utility function is used to generate the welfare weights. A range of non-utility information that may affect how society values the welfare of particular household is not used. The commonly used Atkinson utility functions, for example, rely solely on the resource of the household, so that the underlying assumption of Utilitarianism is that society cares most about the poorest household.

A number of alternative social objectives have been discussed in the the context of optimal tax design and reform. Weinzierl (2013) identifies the principle of Equal Sacrifice as a potential alternative to Utilitarianism, while Saez and Stantcheva (2013) propose a generalised framework in which several objectives, such as Libertarianism and poverty alleviation, can be captured. Similarly, Pirttila and Tuomala (2004) provide a framework whereby poverty alleviation, rather than Utilitarianism, is the objective of the government when designing taxes. In this framework, they show that necessity goods should be taxed lower than other goods, and effective marginal tax rates should be negative for low earners, in order to alleviate poverty.

In the following sections, we turn to identifying the distributions of welfare weights under a number of these alternative approaches. Moving beyond the pure Utilitarianism framework, Weinzierl (2013) showed that some normative frameworks incorporate concerns such as “freedom, rights, and rules, rather than the ends emphasised by Utilitarianism”. Saez and Stantcheva suggested that alternative social preferences can be captured within a generalised social welfare weight framework, where such information can be used. While Weinzierl and Saez, as well as Kuziemko et al. (2013), use specifically designed online surveys to identify plausible alternatives to Utilitarianism, in the following sections, we examine how a number of the proposed alternatives to Utilitarianism can be incorporated using a Household Budget Survey from Ireland. Under each specification, we examine how the distribution of weights differs from the traditional Utilitarian approach.

4.3 Equal Sacrifice and the Libertarian Criterion

Due to the limitations of the Utilitarian approach discussed above, a growing literature suggests alternative determinants of the profile of social marginal welfare weights. Two such alternatives are the related concepts of Equal Sacrifice and Libertarianism. Weinzierl (2013) questioned the “default assumption” of Utilitarianism in tax design and tax reform models. Using an online survey, Weinzierl asked respondents’ opinions on a range of hypothetical tax and benefit systems. The systems ranged from one based on pure Utilitarianism, where the poorest households had a negative average tax rate subsidised by high taxes on the rich households, to one based on the principle of Equal Sacrifice, where each household faced the same

¹⁷While Saez and Stantcheva identify this as a criticism of the Utilitarian approach, it could equally be seen as a criticism of the zero elasticity of labour assumption.

¹⁸Examples of tags include height, race or gender. These are observable characteristics which are correlated with intrinsic earning ability, but are generally deemed unfair to use to determine tax liability

average tax rate. He showed that the vast majority of respondents would choose a system based at least in part on the principle of Equal Sacrifice. In his model, Weinzierl’s approach was to represent each normative criterion with a loss function that depends on deviations in the allocation of resources from the criterion’s optimal allocation. This is an alternative, although related, approach to the social marginal welfare weight approach taken in this paper, and in Saez and Stantcheva (2013). The interrelated nature of the two approaches is evident from Weinzierl’s suggestion that one could translate a mixed objective function from his approach into a set of social marginal welfare weights.

A related concept is that of Libertarianism. Libertarianism in its absolute form, according to Murphy and Nagel (2002), is that no compulsory taxation is legitimate, as the confiscation of income is unjust. Murphy and Nagel argued that “if (and only if) [Libertarianism] is the theory of distributive justice we accept, the principle of Equal Sacrifice does make sense” (Weinzierl 2013).

The principles of Equal Sacrifice and Libertarianism result in a similar motivation for determining the profile of social marginal welfare weights. In both cases, the highest weight will be placed on the household that pays the highest proportion of its gross income in tax payments, while the lowest weight will be placed on the household that pays least tax and receives most state transfers¹⁹. The principle of Equal Sacrifice suggests that each household should face the same average tax rate, so that those households facing a relatively high average tax rate should have the highest weight placed on their welfare in any tax reform scenarios. The extreme Libertarian view is that the confiscation of income is unjust, so that households losing the highest proportion of income through the tax system should be the first ones compensated by tax reform. As Saez and Stantcheva (2013) argue, the Libertarian view is that the level of one’s disposable income is irrelevant and only the tax contribution matters for how socially deserving an individual is. With social marginal welfare weights motivated in this way, a progressive transfer may lead to a fall in social welfare.

We can specify a simple function to model the Equal Sacrifice or Libertarian welfare weights, which indicate that the higher the average tax rate faced by a household, the more “deserving” it is in tax reform scenarios. In contrast to the Utilitarian case where welfare weights are decreasing in a given parameter, here welfare weights are increasing in the parameter.

$$\beta^h = \left(\frac{t^h}{t^N} \right)^\phi \quad (4)$$

where t^N is the highest proportion of gross income paid in tax observed in the data. The household with the highest net (equivalised) tax payment will therefore have the highest welfare weight, normalised to 1. By setting $\phi = 1$, the welfare weights motivated by the principle of equal sacrifice increase at a constant rate in the proportion of income paid in tax²⁰.

¹⁹One could also include the value of public goods received, most likely by including it in the value of transfers received. However, the HBS data used in this paper does not allow the consumption of such public goods to be measured

²⁰Our preferred specification here is $\phi = 1$ as it allows constant returns to increasing tax payments. We could not find evidence in the Equal Sacrifice or Libertarianism literature to suggest a higher value on a household moving from paying 10 per cent tax to 11 per cent tax, compared to a household moving from 40 per cent tax to 41 per cent tax. However, one can also impose the desired degree of concavity or convexity to the concentration curve by setting ϕ not equal to one in equation 4

Again, we normalise the welfare weights so that the household with the highest proportion of income paid in tax has the highest welfare weight, equal to one, as per equation 4.

One problem that arises with this approach is that when basing the welfare weights on net tax and benefit payments, we observe negative values for t^h in the cases where a household receives more in state transfer payments than it pays in taxes. This results in negative welfare weights²¹. To get around this issue, in cases where we observe negative t^h , we normalise the welfare weight by:

$$\beta^{*h} = \left(\frac{t^h - t^1}{t^N - t^1} \right)^\phi \quad (5)$$

where t^1 is the minimum observed value of t^h .

We are faced with a number of choices about the measure to use for t^h in this case. The first measure we use, $w(tt)$, is a measure of the net payment between household and state. It includes the direct and indirect tax payments made by the household, minus the state transfer payments the household receives. It therefore acts as a measure of the net payment between household and state²². Although our data allows us to at least estimate both direct and indirect tax payments made by each household, as well as state transfer payments received, in some cases researchers may be restricted to using just direct taxes and benefits, or indirect tax payments. In order to test the sensitivity of our measure to the exclusion of either tax system, we construct two alternative measures of t^h . The first alternative is to set t^h equal to the direct taxes paid by a household minus the benefits received by a household, expressed as a proportion of gross income, which results in a set of welfare weights $w(dt)$. If the researcher is restricted by data constraints so that he can only measure indirect taxes paid by a household, an alternative might be to set t^h equal to indirect tax paid as a proportion of gross income, resulting in $w(it)$. Household expenditure microdata alone, for example, allows the researcher to approximate the amount of indirect tax paid by a household, but not the direct taxes paid.

Figure 2 shows the distribution of welfare weights based on these three measures of net tax paid. The concentration curve of $w(dt)$ has a high degree of convexity in income, indicating that the welfare weights are strongly concentrated in the top income deciles. Interestingly, the weights are not monotonically increasing in income - the bottom decile has a larger share of weights than the second decile, resulting in the relatively steep section of the curve at the bottom of the income distribution. Callan et al. (2013) show that over the 2008 to 2011 period in Ireland, increased welfare payments offset falls in employment income for those in the 2nd and 3rd deciles, but falls in employment income in the bottom decile were not offset by increases in welfare payments. A similar pattern emerges here, whereby those at the very bottom of the income distribution receive a smaller share of their income in transfer payments than those slightly higher up the income distribution²³. Despite having lower income, those

²¹Non-negative welfare weights ensure any tax optimum is constrained Pareto efficient. See Saez and Stantcheva (2013)

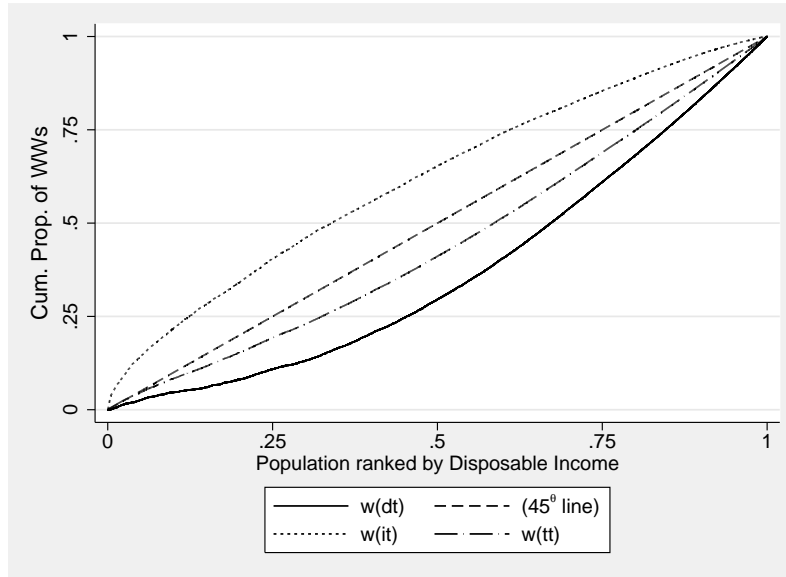
²²An extension of this measure might be to include the value of public goods consumed. However, it is unclear if the consumption of public goods should be included here. With pure public goods, one household's consumption will not affect another household's consumption. This distinction is not as clear with semi-public goods or non-marketed services provided by the state, such as free education. Either way, the HBS data used in this paper does not identify public good consumption so we do not include it in our measure of t^h and leave analysis of this issue to further research.

²³The majority of households in the bottom deciles will not have income high enough to be liable for direct taxation, so patterns of welfare receipt explain the majority of differences for these deciles.

in the bottom decile may have lower welfare payments due to reduced eligibility²⁴, a lack of take-up, or non-eligibility for non-means-tested benefits such as the old age pension or child benefit, for example²⁵.

The distribution of $w(it)$ is quite different than it's direct tax and benefits counterpart. In this case, the weights are decreasing in income, so the concentration curve is concave. This reflects the regressive nature of indirect taxes. Indeed, using a the 1995 wave of the same survey data used in this paper, O'Donoghue and Baldini (2004) showed that the indirect tax system added more than 10 per cent to the Gini Coefficient in Ireland.

Figure 2: Concentration Curve for Libertarian Welfare Weights



The concentration curve of $w(tt)$, the combination of the direct and indirect taxes paid, and benefits received, as a proportion of gross income, dominates both $w(dt)$ and $w(it)$. The low concentration of welfare weights in the low income households in $w(dt)$ is compensated by the high concentration for the same households in $w(it)$. The distribution is convex in income as households in Ireland tend to pay more of their income in direct taxes than indirect taxes.

The approach in this section contrasts strongly with the Utilitarian approach. While the Utilitarian welfare weights presented in the previous section resulted in weights decreasing in income, the Equal Sacrifice/Libertarian weights, in particular when we included direct tax and benefits, results in the opposite. Of course, a distribution of welfare weights similar to the Equal Sacrifice/Libertarian case can be achieved in the Utilitarian framework by specifying a convex transformation of the utility function in the social welfare function, which undoes the concavity of the utility function. However, Saez and Stantcheva (2013) argue that this seems much more artificial than directly stating that society considers redistribution as unjust confiscation.

²⁴For example, previously self-employed individuals may have reduced eligibility to unemployment supports due to a lack of social insurance contributions

²⁵Based on the findings of Brewer and O'Dea (2012), this phenomenon may also be due to misreporting of income in the bottom decile

4.4 Poverty-Based Social Welfare Weights

A measure often used to describe the number of low income households in society is the poverty rate. The “at-risk-of-poverty” (AROP) rate measures the proportion of persons with equivalised disposable income below 60 per cent of the median income, and is used by the European Commission to measure poverty across the EU. Given this rate is often the focus of public debate, public policy decisions can be made with the AROP group in mind²⁶. In a tax reform scenario, a tax analyst may therefore wish to concentrate the welfare weights on those at or below the poverty line.

Foster et al. (1984) provided a class of widely used poverty measures upon which we can measure poverty:

$$P_\alpha(y_d, z) = 1/h \sum_{h=1}^H \left(\frac{g^h}{z} \right)^\alpha \quad (6)$$

where z is the predetermined poverty line, and g_h is the gap from z to y_d^h for households below the poverty line²⁷. The parameter α captures the desired level of poverty aversion. By setting $\alpha = 0$, we have the headcount ratio. Increasing α to 1 we have the poverty gap measure. With $\alpha = 2$ we have the squared poverty gap measure. In general, a larger α places more weight on the poorest poor.

From equation 6, we can construct a poverty-based household weight, by:

$$P_\alpha^h(y_d, z) = \left(\frac{g^h}{z} \right)^\alpha \quad (7)$$

The profile of these weights, which concentrate the weights on the households in poverty, is dependent on the values of the parameters in equation 7. We represent the AROP measure by setting z equal to 60 per cent of median income. We therefore calculate three sets of welfare weights which represent the FGT index with $\alpha = 0$, $w_{\alpha 0}(y_d, z)$, $\alpha = 1$, $w_{\alpha 1}(y_d, z)$, and $\alpha = 2$, $w_{\alpha 2}(y_d, z)$. In all cases, households above the poverty line have a weight of zero. $w_{\alpha 0}(y_d, z)$ is a set of binary weights where AROP households have $\beta^h = 1$. $w_{\alpha 1}(y_d, z)$ and $w_{\alpha 2}(y_d, z)$ have weights which increase in the poverty gap, according to equation 7. $w_{\alpha 1}(y_d, z)$ has the welfare weight given by the gap to the poverty line g_h . $w_{\alpha 1}(y_d, z)$ sets the weight as the square of the gap. In each of $w_{\alpha 1}(y_d, z)$ and $w_{\alpha 2}(y_d, z)$, we normalise the weights so that the household with the highest observed value of g^h has the highest $\beta = 1$.

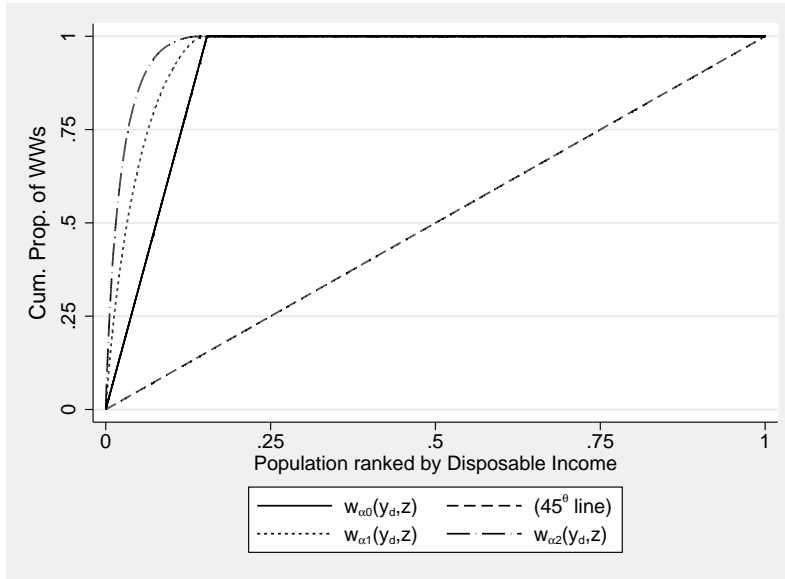
By design, the distribution of weights is fully concentrated in the low income households²⁸. In each case, the concentration curve becomes precisely flat at the poverty line, as households above this line have a zero weight. Figure 3 shows that as α increases, the welfare weights become more and more concentrated in the poorest households. By splitting the population into deciles, this pattern becomes particularly clear. The bottom decile has a 62 per cent share of weights under $w_{\alpha 0}(y_d, z)$, with the remainder of the weights located in decile 2. $w_{\alpha 1}(y_d, z)$ has 89 per cent of weights in the bottom decile, while $w_{\alpha 2}(y_d, z)$ has 98 per cent of weights in the bottom decile. As expected, the higher the level of poverty aversion, the higher the relative weights placed on the poorest household.

²⁶see, for example, the Department of Social Protection’s Social Impact Assessment of Budget 2013, which includes the impact of Budget 2013 on the AROP rate (Department of Social Protection 2013)

²⁷ $g_h = 0$ otherwise

²⁸ z , measured as 60 per cent of median disposable income, is located in the second decile.

Figure 3: Concentration Curve for Poverty-based Welfare Weights



4.5 “Deserved” Income

To this point, we have treated all income in a household equally, regardless of the source or nature of that income. In this section, we question that assumption. Boadway and Martineau (2013) discussed the role of social norms in determining welfare weights. They allow social norms to affect participants and non-participants in the labour market in different ways. In the case when outcomes are determined by luck, preferences for redistribution are increased due the moral benefit of not stigmatizing non-participants for their fate, and redistributing income towards them. On the other hand, when economic outcomes are attributed to effort, the social norm is included as the moral cost for non-participants: this reflects the stigma incurred by those who do not work, as a result of societal attitudes towards idleness. They argue that in the case of a norm embodying the view that outcomes are dictated by luck, the weights on all working classes will be lower, for a given system of taxes and transfers. The opposite result holds when the norm is based on the view that outcomes are the product of effort. The norm where outcomes are due to effort reflects Saez and Stantcheva’s argument that an important belief society seems to hold is that it is fairer to tax income due to ‘luck’ than income earned through hard effort. Using the household budget survey data, we apply two separate approaches in an attempt to identify the element of income that was earned through ‘luck’, and the element that was earned by effort. The aim is to produce a profile of welfare weights where “deserved” or “earned” income is valued higher than “luck” income, so that society more readily accepts the redistribution of the luck element of income.

Our first approach is to base the welfare weights not only on the level of income within a household, but also on the number of hours worked to earn that income²⁹. More formally, we have a set of weights $w(Y, H)$, where the weights increase in the numbers of hours worked, H ,

²⁹Here we work only with households whose head is of working age. Including retirement age households complicates the use of hours of work

and decrease (as before) in income, Y . We have:

$$w(Y, H) = w(Y) * w(H) \quad (8)$$

where $w(Y)$ is as before. $w(H)$ is calculated by $\beta^h = (H^h/H^N)^\phi$ where H_h is the average number of hours worked per person in household h .

This approach captures the view that the more effort (measured here by hours of work) a household exerts to earn its income, the more society cares about the welfare of that household. It also reflects the concept of “full income”, whereby the amount of leisure time available to a household (which reduces in H), as well as disposable income, should also be taken into account. For example, if two households have the same income, but one household works twice as hard as the other to earn that income, the harder working household will have the higher welfare weight, as it has exerted more effort to earn the same level of income, and as a result will have a lower amount of leisure time available. Alesina and Angeletos (2005) suggest that regional variation exists in societies’ beliefs regarding what constitutes ‘luck’ and ‘deserved’ income. Americans, they argue, believe that poverty is due to bad choices or lack of effort while Europeans view poverty as a trap from which it is hard to escape. The approach used here, in which the welfare weights increase in effort but decrease in income, can be seen as a combination of the two. Saez and Stantcheva (2013) find mixed evidence for differing welfare weights based on hours of work. In an online survey conducted by the authors, 42.7 per cent of respondents support a tax break for an individual who works longer hours than another with equal income. 54.4 per cent of respondents, however, think hours of work is irrelevant. Specifying the marginal welfare weights in such a manner may simply be punishing the higher educated by assigning them a lower weight. If hourly wages are a function of education, then two households that have the same (weekly or annual) income but work a different number of hours will have different weights. The higher educated household, which will have a higher hourly wage under the simplified assumptions, would have a lower welfare weight. Careful consideration of the effect on the incentive to invest in education would be necessary under this specification.

The second approach ignores hours of work in determining the element of income that has been earned through ‘luck’, but rather focuses on the source of the income. The HBS data allows us to identify the proportion of income that has been earned through returns on investments or earned through property (which we assign as ‘luck’ income). The design of the tax system in Ireland seems to suggest that society accepts redistribution of the ‘luck’ element of income defined in this way more readily than the residual ‘deserved’ element. While the top marginal income tax rate in Ireland is 51 per cent (including social insurance contributions and the universal social charge), Capital Gains Tax³⁰ reaches as high as 80 per cent³¹. In this setting, we calculate a set of welfare weights $w(Y_{des})$ by equation 3, where I^h is the deserved element of income of household h .

³⁰which is a tax paid upon the sale of an asset such as property


³¹This is quite a simplified example as there are likely to be a range of other factors affecting individual tax rates

Table 2: Luck income as percentage of gross income by decile

Decile	“Luck” Income as % of Gross Income
Bottom	2.0
2	0.5
3	0.8
4	1.6
5	1.3
6	1.0
7	1.2
8	1.6
9	1.2
Top	2.5

Figure 4 shows the distribution of welfare weights based on these two approaches. The concentration curve of the standard Utilitarian approach using gross income as the measure of resource is also replicated for reference. The distribution of $w(Y_{des})$ is remarkably similar to that of $w(Y)$. As can be seen in Table 2, the distribution of “luck” income is concentrated in the top and bottom deciles. The top decile has the highest proportion of income made up of “luck” income, yet it only makes up 2.5 per cent of income in that decile. The second and third deciles have the lowest proportion of income consisting of “luck” income, with each between 0.5 and 0.8 per cent. Consequently, deciles 2 and 3 have a slightly larger share of $w(Y_{des})$ than $w(Y)$, yet the differences are small given the low share of “luck” income in each decile.

Figure 4: Decile Share of “Deserved” Income Welfare Weights



The distribution of $w(Y, H)$ is interesting in a number of respects. The concentration curve is above the 45 degree line for the bottom and top of the income distribution, yet is below the 45 degree line for the middle of the income distribution. Therefore, the welfare weights are increasing in income in the middle of the distribution, yet decline in income at the extremes of the income distribution. This suggests that for the middle of the income distribution the effect of the increasing number of hours worked is stronger than the increasing income. In the extremes of the distribution, the opposite is true, where the income effect is stronger than the hours effect. On average, $w(Y, H)$ is similar to the extreme Utilitarian case with $e = 0$, with a concentration index of close to zero (see Table 7 in the appendix). From Table 3 we can see why this pattern emerges. Hours of work increases strongly from decile 3 to 7, while income growth is less pronounced, so the concentration curve is convex for the middle of the income distribution. On average, those in the bottom decile work longer and earn less than those in the 2nd decile. In the top three deciles, the increase in income growth is larger than the increase in work hours, so the concentration curve again becomes concave.

Table 3: Mean Hours Work and Disposable Income by Decile, for households with working age head

Decile	Mean Hours	% Change on previous decile	Mean Disposable Income	% Change previous decile
Bottom	5.4	—	159	—
2	4.2	-24	217	36
3	7.7	85	253	17
4	10.6	37	292	15
5	15.2	43	340	16
6	19.1	26	403	18
7	24.0	26	481	19
8	25.7	7	581	21
9	27.7	8	722	24
Top	30.0	9	1,126	56

4.6 Governmental Self-Interest

In the previous four subsections, we examined a range of possibilities for the way in which society might determine which household is most deserving. The Utilitarian approach suggests that the household with the least resources is most deserving, while under the principle of Equal Sacrifice the household that contributes the highest proportion of its income in taxes is considered most deserving. But the question remains: when making tax reform decisions, do policy makers place most emphasis on the welfare of the household that society cares most about, or do they have their own motivations for valuing the welfare of certain households? In this section, we examine such a possibility. A literature exists which suggests that the primary aim of an incumbent government is to get reelected, rather than maximise social welfare. Suiter and O'Malley (2013) argued that political parties will often target distributive spending on particular groups of voters. They argued that, depending on the type of electoral system and decision-making rules in place, such public spending can be important for the re-

election of certain candidates or to help a party achieve or maintain a majority in parliament. This may be particularly true in Ireland, they reasoned, because there is strong incentive to garner the personal vote. Similarly, Kayser (2005) suggested that incumbent governments are more likely to enact expansionary economic policy before election time in order to maximise it's chances of reelection. In such a scenario, a policy maker may not place most weight upon the welfare of the household that society values most, but may use a distribution of weights such that he maximises the chances of reelection³².

We adopt two approaches to model this scenario. First, we follow the median voter literature. As Romer and Rosenthal (1979) reason, in the crudest version of the economic theory, the median voter has median income. We therefore maximise the weight on the household with the median disposable income. This approach gives us a starting point for examining the distribution of weights under a governmental self-interest motivation, and reflects what Piketty and Saez (2013) term “the most popular model for policy decisions among economists”. Formally, we calculate the weights as:

$$\beta^h = \left(\frac{x}{g_{my}^h} \right)^\phi \quad (9)$$

where x is the expenditure of the median voter normalised to one, and g_{my}^h is the absolute value of the difference between income of household h and the median income, with minimum value of one. This specification results in a profile of welfare weights $w(y_{med})$.

Of course, the median voter framework is a considerable simplification of a complex political process, requiring single peaked preferences a with respect to a unidimensional policy decision. Romer and Rosenthal found no evidence that empirical results based on the median voter model are superior to any reasonable alternative models. With this in mind, a number of authors have attempted to identify the group of voters that a government would target in order to maximise it's likelihood of re-election. From a review of the political economy literature, it seems unclear which group of voters this represents. Cox and McCubbins (1986) suggested that politicians will target their core supporters, as these voters promise a higher return on their investment. However, Lindbeck (1987) theorised that expected returns are maximised by targeting swing voters, as core voters' choices are generally between voting for their own party or non-participation. More recently, Suiter and O'Malley (2013) suggested that only once factors such as decision-rules have been taken into account can we assess whether core or swing voters are the primary target of partisan spending. Due to this uncertainty, we construct an alternative model of governmental self-interest whereby the policy maker maximises the welfare weight on those most likely to vote, ignoring the distinction between core and swing voters. To model this scenario, we use datasets which record voter participation to model individual level likelihood to vote.

Ideally, we would like to use voter participation modules from the same time period as the main HBS data. However, such a dataset does not exist in Ireland³³. The closest available data comes from a Quarterly National Household Survey (QNHS) special module on voter participation from the 2002 Irish General Election.

³²For simplicity, we take the government and policy maker as one individual decision-making unit. In reality of course, multiple decision makers exist within most governments, and so each may use a different distribution of welfare weights when making decisions. The electoral and decision making rules may also affect the policy decisions made in government (see Suiter and O'Malley (2013)).

³³This reflects the absence of a major election in 2009/10 rather than a shortcoming of data collectors

We combine the main HBS data with the QNHS voter dataset to construct a household level probability to vote measure using a two-step procedure. In the first step, we estimate a parsimonious individual level model estimating the probability of voting with the QNHS data, based on K characteristics identifiable in both datasets. We have:

$$\pi_i^{qnhs} = Pr(Y = 1 | X_1, X_2, \dots, X_K) = F(\gamma_0^{qnhs} + \sum_{k=1}^K \gamma_k^{qnhs} X_k^{qnhs}) \quad (10)$$

In the second stage, these estimates are used to construct a likelihood to vote in our main dataset, the HBS, so that:

$$\hat{\pi}_h^{hbs} = L(\gamma_0^{qnhs} + \sum_{k=1}^K \gamma_k^{qnhs} X_{kH}^{hbs} + \gamma_0^{qnhs} S_h + \sum_{k=1}^K \gamma_k^{qnhs} X_{kS}^{hbs} S_h) \quad (11)$$

where $L(x)$ returns the inverse logit of x , S_h is a dummy variable indicating the presence of a spouse of the head of household, X_{kH} are the characteristics of the head of household, and X_{kS} are the characteristics of the spouse of head of household for $k = 1$ to K . As the HBS only records the required information (X_1, X_2, \dots, X_K) for the head of household and the spouse, we can only construct a likelihood to vote for two people per household in the HBS. $\hat{\pi}_h^{hbs}$ is therefore a household level probability of voting measure.


The distribution of $\hat{\pi}_h^{hbs}$ is normalised so that the likelihood in the household with the highest combined likelihood to vote is equal to one. We use this normalised household level probability to vote to construct a set of welfare weights, $w(\pi)$.

Of course, it may be that while policy makers place a higher weight on those most likely to vote, they are most interested not only in those most likely to vote, but in those most likely to have an influential vote. In that sense, a final specification in this setting is to capture elements of both $w(y_{med})$ and $w(\pi)$, by simply interacting the two sets of weights. This new set of weights, $w(y_{med}, \pi)$, maximises the welfare weight on households most likely to vote that are closest to the median voter, by allowing a one-for-one trade off between the two factors.

The sets of weights proposed in this section can be seen as questioning the the principle of anonymity. Put simply, as argued by Sethuraman et al. (2006), anonymity requires that the preferences of a bandit receive as much consideration as those of a benedictine. In the previous sections, even though someone might not agree with the criterion used to determine the welfare weights, each set of weights was impersonal. With the poverty based weights for example, if household income was furthest below the poverty line, that household got the highest weight. Here the weights are not impersonal, as policy is applied to maximise welfare of a certain set of voters, in order to maximise re-election chances.

Figure 5 shows the distribution of welfare weights under the three governmental self-interest approaches. $w(\pi)$ is convex in income, although is quite close to the 45 degree line at all points. It therefore exhibits a relatively weak positive relationship with income. The concentration curve for $w(y_{med})$ is steepest in the middle of the income distribution, indicating that this profile of weights are highly concentrated among middle income households. This is unsurprising given the specification of these weights in equation 9. $w(y_{med}, \pi)$ is closely related to $w(y_{med})$, which is unsurprising given the reasonably equal distribution of $w(\pi)$.

Figure 5: Decile Share of Govt. Self Interest Welfare Weights



4.7 Summary

The analysis in this section has examined the implications for the distribution of social marginal welfare weights of moving beyond Utilitarianism. The alternatives to Utilitarianism provided here are far from exhaustive. Even within each section, there is an almost infinite number of specifications that could capture the key characteristics of the preferences discussed. For example, in the poverty section, rather than using the AROP measure of poverty, one could focus on household joblessness which has been argued to be an equally important determinant of poverty risk as income (see Watson et al. 2013). Similarly, one could use a range of other characteristics for what is considered “luck” income and “deserved” income in Section 4.5. In the same section, rather than using the normalised product of $w(Y)$ and $w(H)$ to characterise $w(Y, H)$, one could simply use the sum.

The use of a single cross sectional dataset in this paper means we are unable to identify lifetime earnings or tax payments, rather than the static measures we observe. In the Equal Sacrifice section, for example, households dependent solely on the state pension would have a relatively low welfare weight, yet may have made large tax payments in previous periods of employment. This is also an issue within the Utilitarian framework, as low current-income households (such as those dependent solely on the state pension) will have a relatively high welfare weight, yet may have large savings or assets from a previous period of high earnings.

Much analysis of marginal tax reform has concentrated upon sensitivity to either specification of the underlying consumer demand model and/or differential welfare weights within a specifically Utilitarian framework (see Decoster and Schokkaert (1990), Madden (1995)). In the application which follows, we analyse how tax reform proposals change when we step outside the Utilitarian framework and adopt the alternative approaches to determining welfare weights outlined above.

5 Marginal Indirect Tax Reform

5.1 Distributional Characteristic

In this section we use a simple indirect tax reform model to test the sensitivity of results to changes in the underlying social marginal welfare weights. As we want to focus entirely on distributional considerations rather than any efficiency or other concerns that may also influence indirect tax rates (thereby isolating the effect of changes in the welfare weights), our tax reform “recommendations” are based purely on the distributional characteristic for each good. The distributional characteristic of a good, which originates in Feldstein (1972), provides a measure of the degree of concentration of consumption of that good within a particular group of households. In an analysis of distributional characteristics of several goods in Ireland, Madden (2009) states that “the most frequently used approach is to let β^h for each household equal its marginal utility of income, where the utility of income function is the well-known Atkinson one”. Here we move beyond the traditional use of Utilitarianism to test the sensitivity of the results to alternative specifications of the welfare weights.

With household consumption of good i denoted x_i^h and aggregate consumption of that good denoted X_i , the distributional characteristic is calculated by:

$$D_i = \frac{\sum_h \beta^h x_i^h}{X_i} \quad (12)$$

Only the ranking, rather than the absolute value, of the D_i is important here. The ranking of the goods by D_i indicates among which groups consumption of particular goods are concentrated. Higher values of D_i indicate that consumption of good i is concentrated among the “deserving” households, however “deserving” is determined within that particular profile of β^h . Within the Utilitarian framework, for example, goods with a high concentration of expenditure among poorer households will have a relatively high D_i . The distributional characteristic therefore gives the optimal direction of tax reform taking only equity considerations into account, with the recommendation being to reduce the tax on goods with relatively high D_i and increase the tax on goods with a relatively low D_i . The absolute value of D_i will be affected by factors such as normalisation and the value of e chosen for each set of β^h , meaning comparisons of D_i across sets of β^h are not meaningful.

Using the HBS data, we can identify household expenditure across a wide range of goods. For the purpose of this analysis, we aggregate expenditure items into ten aggregate goods³⁴. For each of the goods, we calculate the distributional characteristic by equation 12. Table 4 presents the correlation coefficients from the ranking of the distributional characteristics from each set of welfare weights compared to the standard utilitarian case with disposable income as the measure of resource and with $e = 1$. Within the Utilitarian framework, the rankings are relatively insensitive to the value of inequality aversion e . Despite the increasing concentration of welfare weights on the poorest households as e increases (see Figure 1), the rankings of goods by distributional characteristic remains quite similar, with a correlation coefficient of 0.8 between the rankings when $e = 1$ and $e = 5$.

Comparing the Utilitarian rankings with the non-Utilitarian rankings, the most contrasting rankings are between the Utilitarian and Equal Sacrifice or Libertarian weights. The

³⁴The ten goods are food, alcohol, tobacco, clothing, fuel, household non-durables, household durables, housing, transport, and services and other goods

correlation between the ranking of goods when D_i is based on $w(tt)$ and $w(y_d)$ is low, at -0.2. The majority of the contrast in the rankings is driven by the direct tax payments. A correlation coefficient of -0.9 between the rankings of goods based on $w(dt)$ and $w(y_d)$ indicates that the rankings of the goods, and so the tax reform recommendations, are almost the direct opposite between the Utilitarian case and the Equal Sacrifice case based on the direct tax and benefit system. Including the indirect tax payments in the Equal Sacrifice profiles of weights increases the degree of correlation between the rankings. $w(it)$ produces a ranking of goods with a positive correlation of 0.4 with the Utilitarian ranking.

The poverty based welfare weights produce similar distributional characteristic rankings as $w(y_d)$, with the correlation coefficient as high as 0.9 when $\alpha = 0$. As α becomes larger, the rankings diverge³⁵. This is unsurprising given the poverty based weights are effectively a restricted case of the Utilitarian weights.

While $w(y_{des})$ and $w(y_d)$ produce almost identical rankings of distributional characteristics, $w(y, H)$ produces quite different rankings, with a correlation coefficient of -0.5. Interestingly, while $w(\pi)$ is concave in income, the resultant rankings of D_i contrast with the Utilitarian rankings, with a correlation coefficient of -0.4.

Table 4: Correlation Coefficients for Ranking of Goods by Distributional Characteristic

	$w(y_d)$		$w(y_d)$		$w(y_d)$
$w(y_d^{e=2})$	0.99	$w_{\alpha 0}(y_d, z)$	0.89	$w(y, H)$	-0.54
$w(y_d^{e=5})$	0.82	$w_{\alpha 1}(y_d, z)$	0.73	$w(Y_{med})$	0.75
$w(dt)$	-0.95	$w_{\alpha 2}(y_d, z)$	0.47	$w(\pi)$	-0.37
$w(it)$	0.43	$w(y_{des})$	0.99	$w(Y_{med}, \pi)$	0.87
$w(tt)$	-0.22				

Table 5 shows the ranking of goods by D_i under each profile of welfare weights. As was shown in Table 4, the rankings are sensitive to the choice of welfare weights. In the Utilitarian setting, food, fuel and tobacco are the highest ranked goods, indicating that expenditure on these goods is concentrated among income poor households. Services and housing are at the bottom of the rankings, so that richer households spend relatively higher proportions of their incomes on these goods. The tax reform recommendations based on the Utilitarian rankings, taking only equity considerations into account, would therefore be to reduce the tax on the goods with the higher D_i , such as tobacco and fuel, and increase the tax on the lower ranking goods, such as services and, particularly at higher levels of inequality aversion, alcohol.

These tax reform recommendations change significantly as we change the profile of welfare weights. Under $w(dt)$ and $w(tt)$, services, alcohol and transport are the highest ranking goods. Tobacco is consistently towards the bottom of the rankings under these weights, joined by food and fuel, or housing and durables, depending on the choice of welfare weights. $w(it)$ produces a ranking of goods more similar to the Utilitarian case, although alcohol has a relatively high D_i under $w(it)$, which is unsurprising given the high level of indirect taxation that alcohol is subject to in Ireland³⁶.

³⁵Increasing the value of e in the Utilitarian weights maintains the high correlation coefficients as α increases in the poverty based weights.

³⁶As well as being subject to the standard rate of VAT, it is also subject to excise duties.

The poverty-based weights produce a ranking of D_i close to the Utilitarian rankings. Fuel and transport are most sensitive to the value of α , with fuel getting a relatively lower D_i as α increases. Alcohol and tobacco are the bottom of the rankings under $w(y, H)$, while they are at opposite ends of the rankings with $w(y_{med})$.

Table 5: Rankings of Goods by Distributional Characteristic - Alternative Welfare Weights

	$w(y_d^{e=1})$	$w(y_d^{e=2})$	$w(y_d^{e=5})$	$w(dt)$	$w(it)$	$w(tt)$	$w_{\alpha 0}(y_d, z)$
Alcohol	7	7	10	3	2	2	5
Clothing	5	6	6	6	4	4	8
Food	3	3	4	8	8	6	3
Fuel	2	2	1	9	6	7	2
Non Durables	4	4	2	7	7	5	4
Services	10	10	7	1	10	1	10
Tobacco	1	1	3	10	1	8	1
Transport	6	5	5	4	3	3	6
Durables	8	8	8	5	5	9	9
Housing	9	9	9	2	9	10	7

	$w_{\alpha 1}(y_d, z)$	$w_{\alpha 2}(y_d, z)$	$w(y_{des})$	$w(y, H)$	$w(Y_{med})$	$w(\pi)$	$w(Y_{med}, \pi)$
Alcohol	3	3	7	8	9	8	9
Clothing	5	4	6	5	8	7	7
Food	4	5	3	7	5	4	3
Fuel	2	7	2	9	3	6	4
Non Durables	8	8	4	3	2	5	2
Services	10	10	10	4	10	1	10
Tobacco	1	1	1	10	1	10	1
Transport	7	2	5	1	4	2	5
Durables	9	9	8	2	6	3	6
Housing	6	6	9	6	7	9	8

Note: Goods ranked by distributional characteristic - Highest(1) to lowest (10)

The rankings in Table 5, and the related correlations in 4, show that indirect tax reform recommendations can be quite sensitive to the specification of the underlying social marginal welfare weights. Within the Utilitarian framework, changing the value of e has a relatively small impact on the tax reform recommendations that result from the distributional characteristic analysis. Moving beyond the Utilitarian framework, it is only when we include non-income information that we observe significant changes in the rankings of goods. Allowing the welfare weights to increase in tax payments, in hours of work, or in the probability of voting, all produce rankings of goods with negative correlations to the traditional Utilitarian ranking. The poverty based weights, which are functions of income alone, produce rankings with high correlations to the Utilitarian case.

Of course, the simple tax reform model used here ignores many important considerations that would be taken into account when designing or reforming a tax system, such as efficiency

and externality-correcting concerns. Adding such considerations to the tax reform model³⁷ would likely reduce the sensitivity of results to changes in the welfare weights. Nonetheless, the use of the distributional characteristic has highlighted that moving beyond Utilitarianism in determining the marginal social welfare weights can significantly alter the impact of equity concerns in tax reform modelling. As stated previously, the aim of this paper is not to argue in favour of a particular set of weights. Rather, it is hoped that the analysis in this paper makes clear the implications of such a choice; one which is often made *ex ante* by the researcher.

6 Conclusion

Social marginal welfare weights play an important role in areas of applied public policy analysis such as tax reform. These weights reflect the values of the social planner, or equivalently the underlying social welfare function. Thus when evaluating the welfare effects of a tax reform, different weights may be assigned to different groups in society, with the weights reflecting the relative importance of these groups in terms of aggregating gains and losses arising from the tax reform. Until recently, the majority of research has adopted social marginal welfare weights based upon a Utilitarian approach, with concave utility of income functions, so that the poorest agent (individual or household) in society receives the highest welfare weight and weights decline as income increases. The rate at which the weights decline in income is determined by the level of inequality aversion specified in the social welfare function.

A number of recent contributions to the literature have questioned the default use of welfarism in this setting (Saez and Stantcheva (2013), Weinzierl (2013)). Using the nationally representative 2009/10 Irish Household Budget Survey, we apply a range of alternatives to welfarism in determining the distribution of social marginal welfare weights, and compare these distributions to that arising from the traditional welfarist approach. With these alternative distributions of welfare weights, we then examine whether optimal tax reform recommendations are sensitive to a move away from welfarism.

The distribution of welfare weights arising from the alternative approaches are found to differ appreciably from the distribution based upon welfarist weights. In applying alternatives to welfarism, weights are based on a positive or negative transformation of income (poverty, median voter), on income along with some other factor or household characteristic (“deserved” income, equal sacrifice), or on some set of characteristics independent of income (propensity to vote). The alternative weights which rely solely on some monotonic transformation of income produce a distribution of weights that is either very similar to the welfarist distribution, or almost a mirror image of the welfarist distribution, depending on whether the transformation of income is positive or negative. When factors other than income are included in the determination of the marginal social welfare weights, the comparison with the welfarist case becomes more varied. Including these other factors often produces distributions of weights that are not monotonic in income, and so the implications for tax reform becomes less predictable.

To test the sensitivity of indirect tax reform recommendations to changes in the social marginal welfare weights, we calculated the distributional characteristic of a range of goods under each of the sets of weights. The distributional characteristic allowed us to focus on equity considerations in making tax reform recommendations, abstracting from efficiency and other considerations that may also be taken into account in full indirect tax reform models.

³⁷See Ahmad and Stern (1984) for an example of a tax reform model that allows for equity and efficiency concerns.

The degree of correlation between the tax reform recommendations using the welfarist weights and the alternative weights varied from close to perfect correlation (with the poverty-based weights), to a highly negative correlation (with the equal-sacrifice weights). This highlights the importance of careful consideration of the profile of marginal social welfare weights in making tax reform recommendations, as the recommendations, based on equity considerations at least, can be highly dependent on how the weights are determined.

The various alternatives to welfarism examined in this paper are by no means exhaustive. We are restricted in the factors we can include in the analysis by the data available. Saez and Stantcheva (2013), for example, suggest a distinction between “deserving transfer beneficiaries” vs “free loaders” in determining the profile of social marginal welfare weights, where “free loaders” are those who choose to claim benefits rather than work. Using a cross-sectional household budget survey, it is difficult to distinguish between “deserving transfer beneficiaries” and “free loaders” without making some strong assumptions. Further work on this topic that uses alternative data sources, such as panel or administrative data, may be able to include factors that were not possible here.

Given the importance of social marginal welfare weights in areas of public policy analysis such as optimal labour and commodity tax design, and tax reform evaluation, we believe this detailed examination of the alternatives to welfarism, and their application to a household budget survey dataset, is an important addition to the literature. This is particularly true following the studies of Saez and Stantcheva, and Weinzierl, among others, which provide empirical evidence that people’s views on tax reform/public policy issues often seem to deviate from welfarist principles. While these studies provided evidence of a demand for social welfare weights based on principles other than welfarism, to the best of our knowledge this is one of the first attempts to empirically examine the sensitivity of actual tax reforms to a wider range of welfare weights than those implied by welfarism.

A Logit Results for Probability to Vote

Table 6 shows the characteristics associated with a higher probability of voting. Those in the largest urban area in the country, Dublin, are considerably less likely to vote than those elsewhere. Unemployed individuals are less likely to vote than employees, while married individuals are more likely to vote than all other marital statuses. The probability of voting increases with age and education. Usual hours of work seems to have no effect on the voting participation decision, conditional on the other variables.

B Concentration Indices

Table 7 shows the concentration indices for each of the distributions of welfare weights across the income distribution. Graphically, these can be viewed as a measure of the area between the 45 degree line and the concentration curve for each set of welfare weights. A negative value indicates that the weights are more concentrated in low income households, while a positive value indicates that the weights are more concentrated in high income households. A larger value for the concentration index, in absolute terms, indicates a more unequal distribution of weights across the income distribution. The dominance of one curve over another can be inferred by the absolute size of the coefficient, where a smaller value dominates a larger value.

Table 6: Logit coefficients for likelihood to vote - QNHS 2002

Variable	Category	Coefficient	Standard Error
Sex	Female	0.036	(0.043)
Region	SW, SE, MW, ME	-0.146***	(0.042)
	Dublin	-0.362***	(0.044)
PES	Employee (PT)	-0.013	(0.074)
	Assisting Relative	0.633	(0.388)
	Self-Employed (FT)	0.163**	(0.083)
	Self-Employed (PT)	-0.325*	(0.168)
	Unemployed (Seeking Work)	-0.565***	(0.117)
	Unemployed (Other)	-0.364**	(0.141)
	Home Duties	-0.225**	(0.090)
	Retired	-0.159	(0.107)
	In Education	-0.060	(0.143)
	Other	-0.545***	(0.120)
Marital Status	Married	0.602***	(0.070)
	Separated	-0.272***	(0.075)
	Widowed	-0.005	(0.075)
Age	25-34	0.470***	(0.068)
	35-44	1.177***	(0.075)
	45-54	1.707***	(0.081)
	55-64	1.988***	(0.089)
	65+	2.073***	(0.100)
Education	Secondary	0.405***	(0.049)
	Technical or Vocational	0.466***	(0.069)
	Third Level	0.570***	(0.059)
	Other	0.367**	(0.167)
Usual Hours	UsualHrs	-0.000	(0.002)
Farmers	Farmers	0.502***	(0.126)
Household Composition	Lone Parent	0.258***	(0.081)
	Couple with Children	0.283***	(0.055)
	Other	0.213***	(0.071)
Constant	Constant	-0.427***	(0.123)
Observations		24,805	
Standard errors in parentheses			
*** p<0.01, ** p<0.05, * p<0.1			

A complication arises when the concentration curve of interest crosses the 45 degree line, as the positive element of the concentration index (section of curve below the 45 degree line) will be “cancelled out” by the negative element (section of curve above the 45 degree line). A distribution of weights fully concentrated on the median income household would therefore have the same concentration index as an equally distributed set of weights. These indices should therefore be interpreted in combination with the concentration curves.

Table 7: Concentration Indices for Social Marginal Welfare Weights across Income Distribution

Welfare Weight	Concentration Index
$w(E)$	-0.20
$w(y)$	-0.34
$w(y_d)$	-0.31
$w(dt)$	0.28
$w(it)$	-0.23
$w(tt)$	0.12
$w_{\alpha 0}(y_d, z)$	-0.85
$w_{\alpha 1}(y_d, z)$	-0.92
$w_{\alpha 2}(y_d, z)$	-0.96
$w(y, H)$	0.03
$w(Y_{des})$	-0.34
$w(E_{med})$	-0.08
$w(\pi)$	0.03
$w(E_{med}, \pi)$	-0.05

References

- Ahmad, E. and N. Stern (1984). The theory of reform and indian indirect taxes. *Journal of Public Economics* 25(3), 259 – 298.
- Alesina, A. and G.-M. Angeletos (2005, September). Fairness and redistribution. *American Economic Review* 95(4), 960–980.
- Atkinson, A. B. (1970, September). On the measurement of inequality. *Journal of Economic Theory* 2(3), 244–263.
- Boadway, R. and N. G. Martineau (2013). Normative analysis with societal constraints. CESifo Working Paper Series 4305, CESifo Group Munich.
- Brewer, M. and C. O’Dea (2012, July). Measuring living standards with income and consumption: evidence from the uk. IFS Working Papers W12/12, Institute for Fiscal Studies.
- Callan, T., B. Nolan, C. Keane, M. Savage, and J. R. Walsh (2013, June). Crisis, response and distributional impact: The case of ireland. IZA Discussion Papers 7481, Institute for the Study of Labor (IZA).
- Collins, M. L. and D. Turnbull (2013). Estimating the direct and indirect tax contributions of households in ireland. Working Paper 8, Nevin Economic Research Institute.
- Coulter, F. A. E., F. A. Cowell, and S. P. Jenkins (1992). Equivalence scale relativities and the extent of inequality and poverty. *The Economic Journal* 102(414), pp. 1067–1082.
- Cowell, F. A. and K. Gardiner (1999, August). Welfare weights. Technical report, STICERD, London School of Economics. Paper Commissioned by the Office of Fair Trading.
- Cox, G. W. and M. D. McCubbins (1986). Electoral politics as a redistributive game. *The Journal of Politics* 48, 370389.
- Creedy, J. (2006). Evaluating policy: Welfare weights and value judgements. Department of Economics - Working Papers Series 971, The University of Melbourne.
- Dalton, H. (1939). *Principles of Public Finance*. London: Routledge and Kegan Paul.
- Decoster, A. and E. Schokkaert (1990). Tax reform results with different demand systems. *Journal of Public Economics* 41(3), 277 – 296.
- Department of Social Protection (2013, March). Social impact assessment of the main welfare and direct tax measures in budget 2013. Technical report, Department of Social Protection.
- Diamond, P. A. and J. A. Mirrlees (1971, June). Optimal Taxation and Public Production II: Tax Rules. *American Economic Review* 61(3), 261–78.
- Donaldson, D. and J. A. Weymark (1980). A single-parameter generalization of the gini indices of inequality. *Journal of Economic Theory* 22(1), 67 – 86.
- Feldstein, M. S. (1972, March). Distributional equity and the optimal structure of public prices. *American Economic Review* 62(1), 32–36.

- Fleurbaey, M. and F. Maniquet (2007, 06). Help the low skilled or let the hardworking thrive? a study of fairness in optimal income taxation. *Journal of Public Economic Theory* 9(3), 467–500.
- Foster, J., J. Greer, and E. Thorbecke (1984, May). A class of decomposable poverty measures. *Econometrica* 52(3), 761–66.
- H.M. Treasury (2011). Appraisal and evaluation in central government (the green book). Technical report, HM Treasury.
- Kayser, M. A. (2005). Who surfs, who manipulates? the determinants of opportunistic election timing and electorally motivated economic intervention. *American Political Science Review* 99(1), 17–28.
- Kuziemko, I., M. I. Norton, E. Saez, and S. Stantcheva (2013). How elastic are preferences for redistribution? evidence from randomized survey experiments. Working Paper 18865, National Bureau of Economic Research.
- Lindbeck, A. and Weibull, J. (1987). Balanced-budget redistribution as the outcome of political competition. *Public Choice* 52, 273–297.
- Lundberg, S. J., R. A. Pollak, and T. J. Wales (1997). Do husbands and wives pool their resources? evidence from the united kingdom child benefit. *Journal of Human Resources* 32(3), 463–480.
- Madden, D. (1995). An analysis of indirect tax reform in ireland in the 1980s. *Fiscal Studies* 16(1), 18–37.
- Madden, D. (2009, August). Distributional characteristics for ireland: A note. Working Papers 200910, School Of Economics, University College Dublin.
- Mirrlees, J. A. . . (1974). *Essays in Equilibrium Behavior under Uncertainty*, Volume Notes on Welfare Economics, Information and Uncertainty. Amsterdam: North-Holland.
- Mirrlees, J. A. (1971, April). An Exploration in the Theory of Optimum Income Taxation. *Review of Economic Studies* 38(114), 175–208.
- Murphy, L. and T. Nagel (2002). *The Myth of Ownership: Taxes and Justice*. Oxford: Oxford University Press.
- O'Donnell, O., E. van Doorslaer, A. Wagstaff, and M. Lindelow (2008). *Analyzing Health Equity Using Household Survey Data: A Guide to Techniques and Their Implementation*. The International Bank for Reconstruction and Development / The World Bank: Washington D.C.
- O'Donoghue, C. and M. Baldini (2004). Modelling the redistributive impact of indirect taxes in europe: An application of euromod. Working Papers 0077, National University of Ireland Galway, Department of Economics.
- Piketty, T. and E. Saez (2012, November). Optimal labor income taxation. NBER Working Papers 18521, National Bureau of Economic Research, Inc.

- Piketty, T. and E. Saez (2013). *Optimal Labour Income Taxation*, Volume Handbook of Public Economics Vol. 5. North Holland.
- Pirttila, J. and M. Tuomala (2004, October). Poverty alleviation and tax policy. *European Economic Review* 48(5), 1075–1090.
- Ray, R. (1986). Sensitivity of optimal commodity tax rates to alternative demand functional forms: An econometric case study of india. *Journal of Public Economics* 31(2), 253 to 268.
- Revenue Commissioners of Ireland (2012). Statistical report for year ended 31st december, 2011: Excise. Statistical report, Revenue Commissioners.
- Romer, T. and H. Rosenthal (1979). The elusive median voter. *Journal of Public Economics* 12, 143– 170.
- Saez, E. (2001). Using elasticities to derive optimal income tax rates. *The Review of Economic Studies* 68(1), 205–229.
- Saez, E. and S. Stantcheva (2013). Generalized social marginal welfare weights for optimal tax theory. NBER Working Papers 18835, National Bureau of Economic Research, Inc.
- Sethuraman, J., C.-P. Teo, and R. V. Vohra (2006, May). Anonymous monotonic social welfare functions. *Journal of Economic Theory* 128(1), 232–254.
- Suiter, J. and E. O'Malley (2013). Yes, minister: The impact of decision-making rules on geographically targeted particularistic spending. *Parliamentary Affairs*.
- Watson, D., B. Maitre, and S. Cantillon (2013). Implications of income pooling and household decision-making for the measurement of poverty and deprivation: An analysis of the silc 2010 special module for ireland. Technical report, Social Inclusion Technical Paper No. 4 Dublin: Department of Social Protection.
- Weinzierl, M. C. (2012). Why do we redistribute so much but tag so little? the principle of equal sacrifice and optimal taxation. NBER Working Papers 18045, National Bureau of Economic Research, Inc.
- Weinzierl, M. C. (2013). The promise of positive optimal taxation. Working Paper No. 14-013, Harvard Business School Working Paper.
- Young, H. P. (1987, March). Progressive taxation and the equal sacrifice principle. *Journal of Public Economics* 32(2), 203–214.