Food Expenditure and GST in New Zealand

Christopher Ball, John Creedy and Michael Ryan

New Zealand Treasury
Working Paper 14/07

April 2014

THE TREASURY
Kaitohutohu Kaupapa Rawa

New Zealand Government

NZ TREASURY WORKING PAPER 14/07

Food Expenditure and GST in New Zealand

MONTH/YEAR

April 2014

AUTHORS

Christopher Ball
Treasury
No. 1 The Terrace
Wellington
New Zealand
Email: Christopher.Ball@treasury.govt.nz
Telephone: ++64 +4 890 7206

John Creedy
Victoria University of Wellington
23 Lambton Quay
Wellington
New Zealand
Email: John.Creedy@treasury.govt.nz
Telephone: ++64 +4 917 6893

Michael Ryan
Treasury
No. 1 The Terrace
Wellington
New Zealand
Email: Michael.Ryan@treasury.govt.nz
Telephone: No author defined

ISBN (ONLINE)

978-0-478-42162-0

URL

Treasury website at April 2014:
http://www.treasury.govt.nz/publications/research-policy/wp/2014/wp-14-07
Persistent URL: http://purl.oclc.org/nzt/p-1645

ACKNOWLEDGEMENTS

We are grateful to Matt Benge, Angela Mellish, Alastair Thomas and Martin Wong for their comments on an earlier draft of this paper.

NZ TREASURY

New Zealand Treasury
PO Box 3724
Wellington 6008
NEW ZEALAND
Email: information@treasury.govt.nz
Telephone: +64 4 472 2733
Website: www.treasury.govt.nz

Abstract

This paper has two main aims. First, the poor targeting of a policy of zero-rating food in a goods and services tax (GST) is illustrated in a simple model where the revenue lost from zero-rating food is instead devoted to a universal transfer payment, with a larger effect on progressivity. Second, the paper investigates the welfare effects on New Zealand households of zero-rating food. The detailed effects, for a range of household types, are then investigated using Household Economic Survey data. Demand responses to consumer price changes are estimated and welfare changes, in terms of equivalent variations, are obtained. Comparisons are made across 'clusters', consisting of groups of households with similar characteristics. The reform is seen to produce a very small amount of progressivity in the GST, with redistribution from richer households without children to poorer households with children, and older households.

JEL Classification: H31; I3; D11

Keywords: Goods and services tax; tax progressivity; budget shares; welfare changes.

Executive Summary

The New Zealand Goods and Services Tax (GST) is unusual in having a very broad base which includes food. This is usually defended on grounds of efficiency and there is often a presumption in favour of uniformity on the grounds of the large administrative and compliance costs of differentiation. These arguments are combined with the important point that equity objectives can be met using the nonlinear structure of income taxes and transfers.

Nevertheless, other countries use the zero-rating of food (and other goods) for redistributive reasons. The issue is periodically raised in New Zealand, and it is sometimes suggested that if the GST rate becomes much higher than its current level, there is likely to be further pressure, on equity grounds, for the exemption of some goods. Hence, despite the limitations to any analysis of only one part of the complete tax and transfer system, since it is the overall effect that matters, a separate analysis of GST is warranted.

A consumption tax imposed at a uniform rate on *all* goods and services has no redistributive effect since the real value of total expenditure of all households are reduced by the same proportion. A consumption tax is most progressive, or inequality reducing, when it taxes most heavily those goods which form a systematically higher proportion of the budgets of high-expenditure households. This lies behind the argument that food should be exempt, given the long-established empirical relationship of a declining budget share for food as total expenditure increases. But there is a cost of such redistribution because since, to raise the same revenue, the tax rate imposed on other goods must be higher than in the uniform structure. Rational debate requires information both about the costs and benefits of such differentiation.

The present paper contains two analyses of the effects of zero-rating. First, the poor targeting of such a policy is illustrated in general terms in a simple model where the revenue lost from zero-rating food is instead devoted to a universal transfer payment. It is shown that the universal transfer payment has a larger effect on tax progressivity than does zero-rating.

The second analysis computes the detailed effects on New Zealand households of zero-rating food, for a range of household types. The effects were obtained using Household Economic Survey data. Demand responses to consumer price changes were estimated and welfare changes were obtained. Comparisons were made across 'clusters', consisting of groups of households with similar characteristics, as well as for households in total expenditure decile and age groups.

A revenue neutral reform was found to produce a very small amount of progressivity in the GST, with welfare losses for higher decile groups and welfare gains for the lower deciles. Furthermore, redistribution is from richer households without children to poorer households with children, and older households. The value of allowing for demand responses was demonstrated by the fact that the use of tax changes, with an assumption of no responses to price changes, overstates the degree of progressivity achieved.

The analysis thus supports earlier studies suggesting that indirect tax exemptions and zero-rating provide a poor redistributive instrument compared with the use of direct taxes and transfers.

Contents

List of Tables

Food Expenditure and GST in New Zealand

1 Introduction

The New Zealand Goods and Services Tax (GST) is unusual in having a very broad base, which includes food.[1] This is usually defended on grounds of efficiency – that the broad base makes it possible to keep the GST rate relatively low. This keeps the excess burden, which approximately increases with the square of the rate, low. In an optimal tax framework there is no presumption in favour of uniformity; the conditions under which the maximisation of a social welfare function gives rise to uniform indirect taxes are strong and unlikely to hold.[2] However, faced with the enormous difficulty of computing a set of optimal taxes, there is often a presumption in favour of uniformity on the grounds of the large administrative and compliance costs of differentiation.[3] These arguments are combined with the important point that equity objectives can be met using the nonlinear structure of income taxes and transfers.

Nevertheless, other countries have used the zero-rating of food and other goods, along with differential rates, for redistributive reasons.[4] The issue is periodically raised in New Zealand, and it is sometimes suggested that if the GST rate becomes much higher than its current level, there is likely to be further pressure, on equity grounds, for the exemption of some goods. Hence, despite the limitations to any analysis of only one part of the complete tax and transfer system, since it is the overall effect that matters, a separate analysis of GST is warranted. When looking at only one tax, it is important, when measuring tax progressivity, to look at tax

1 For international comparisons, using the concept of ‘C-efficiency’ (reflecting departure from a uniform system), see Keen (2013). He decomposes the efficiency measure into a ‘poverty gap’ and a ‘compliance gap’.

2 This is essentially a second-best problem in view of the impossibility of taxing endowments. For example, with identical preferences, a fixed wage rate distribution, and a linear income tax, uniformity is optimal if Engel curves are linear and the marginal rate of substitution between goods is independent of leisure; see Stern (1990).

3 Indeed, the configuration of tax rates is irrelevant only if judges are indifferent to the distribution of welfare and, in addition, if all own-price and cross-price elasticities of demand are zero. On optimal indirect tax reform in Australia, see Creedy (1999).

4 Reference is often made to the exemption of food from GST. However, such policies impose zero rating, whereby registered firms can claim tax credits on inputs purchased. This differs from the technical term ‘exemption’, where such credits cannot be claimed.

payments in relation to the particular tax base, in this case total expenditure for each household.[5]

A consumption tax imposed at a uniform rate on *all* goods and services has no redistributive effect since the real value of total expenditure of all households are reduced by the same proportion. A consumption tax is most progressive, or inequality reducing, when it taxes most heavily those goods which form a systematically higher proportion of the budgets of high-expenditure households. This lies behind the argument that food should be exempt from a general consumption tax, given the long-established empirical relationship (referred to as Engel's Law) of a declining budget share for food as total expenditure increases. But there is a cost of such redistribution because since, to raise the same revenue, the tax rate imposed on other goods must be higher than in the uniform structure. Rational debate requires information both about the costs and benefits of such differentiation.

The present paper has two main aims. First, the use of differential rates in order to introduce progressivity is discussed in Section 2. Exemptions are compared with the use of the same tax revenue, lost by zero-rating food, to finance a universal benefit. The second aim is to investigate the detailed welfare effects on households of a change to the indirect tax structure in New Zealand involving the zero-rating of food (except for meals consumed outside the home) from GST.[6] A feature of the present paper is that demand responses to consumer price changes are estimated: the approach is discussed in Section 3 and in the Appendix. The treatment of consumers' behaviour makes it possible to estimate money measures of welfare changes arising from tax changes. Comparison of welfare effects are made for a range of household types. In addition, comparisons are made across 'clusters', consisting of groups of households with similar characteristics. The clustering method is also described in Section 3. The welfare changes arising from the current system and a reform involving the zero-rating of food are reported in Section 4. Section 5 briefly concludes.

[5] Spurious results can be obtained if a different tax base, such as income, is used. Popular discussions often focus inappropriately on tax in relation to income. The difference between bases is affected not only by income taxation but, in a short-period framework, by savings. Since savings are ultimately spent, they can be ignored when using annual data. On progressivity and indirect taxes, see detailed treatments in Creedy (1998a, 2001).

[6] For a detailed analysis of zero-rating of food from a legal and administrative point of view, see van Klink (2012).

2 Budget Shares and Tax Structures

This section examines some of the basic considerations to take into account when attempting to use exemptions to achieve distributional objectives.

2.1 Elasticities and Liability Progression

The argument relating to indirect tax progressivity and variations in budget shares with total expenditure can be restated more formally as follows. Let τ_i denote the tax-inclusive *ad valorem* consumption tax rate on good i, for $i = 1, ..., n$ goods. This is related to the tax-exclusive rate, t_i, by $\tau_i = t_i / (1 + t_i)$. The price and consumption of good i are respectively p_i and x_i, and the budget share of i is $w_i = p_i x_i / \sum_{i=1}^{n} p_i x_i = p_i x_i / y$. The total consumption tax, $T(y)$, paid by an individual with total expenditure of y is:

$$T(y) = y \sum_{i=1}^{n} \tau_i w_i \tag{1}$$

One local measure of the progressivity of a tax system is the elasticity of revenue, $T(y)$, with respect to y. The elasticity is also the ratio of the marginal tax rate to the average tax rate. It is in fact the Musgrave-Thin measure of liability progression.[7] It is given by:

$$\eta = 1 + \sum_{i=1}^{n} \left(\frac{\tau_i w_i}{\sum_{i=1}^{n} \tau_i w_i} \right) \left(\frac{dw_i}{dy} \frac{y}{w_i} \right) \tag{2}$$

Let $e_i = 1 + \frac{dw_i}{dy} \frac{y}{w_i}$ denote the total expenditure elasticities, and define $\theta_i = \tau_i w_i / \sum_{i=1}^{n} \tau_i w_i$ as the proportion of tax resulting from good i. then:

$$\eta = 1 + \sum_{i=1}^{n} \theta_i (e_i - 1) \tag{3}$$

Hence η is a tax-share weighted average of the $e_i - 1$.

An increasing budget share is equivalent to $e_i > 1$, and since $\eta > 1$ for a progressive tax structure, (3) indicates that goods with total expenditure elasticities greater than 1 should be taxed most heavily in order to obtain the greatest progressivity. However, some progressivity can be obtained even if taxes are imposed on some goods for which $e_i < 1$ at some total expenditure levels. Furthermore, the elasticities vary with y. Indeed, it is often the case that for goods with $e_i > 1$ at all

7 On the revenue elasticity of indirect taxes in New Zealand, see Creedy and Gemmell (2004).

levels of y, the elasticity declines as y increases.[8] This again provides a constraint on the progressivity of an indirect tax structure, since the tax rate has to be set independently of individuals' total expenditure levels.

Furthermore, there is substantial population heterogeneity. The budget share devoted to food decreases systematically *on average* as total expenditure rises, but there is considerable variation within any total expenditure group. This implies a certain amount of 'reranking' if differential rates are used: the rank ordering of households by after-tax total expenditure is not the same as the ordering by pre-tax expenditure. This reranking reduces the degree of systematic redistribution that can be achieved by zero-rating food.[9]

2.2 A Simple Comparison

Some simple comparisons between alternative structures can be made, assuming for the moment that budget shares are not affected by indirect taxes. Consider a broad-based GST where τ is the common tax-inclusive rate applied to all expenditure. In terms of the tax base, this is thus neutral from a progressivity point of view since the average tax rate, ATR, is τ for all y values. Suppose that food is then zero-rated, but the same rate is applied to all other goods. If the budget share devoted to food declines as total expenditure increases, according to:

$$w = \alpha + \frac{\gamma}{y} \tag{4}$$

The average tax rate, ATR_E, faced by someone with total expenditure of y is thus:

$$ATR_E = \tau(1-w) = \tau\left(1-\alpha-\frac{\gamma}{y}\right) \tag{5}$$

and the consumption tax is progressive because ATR_E clearly increases as y increases. With a population size of P, total tax revenue per person, T_E, is, again ignoring individual subscripts, equal to $T_E = \frac{\tau}{P}\sum y(1-w)$. Hence:

$$T_E = \tau\left\{\bar{y} - (\gamma + \alpha\bar{y})\right\} \tag{6}$$

Instead of zero-rating food, suppose that the tax raised from food is used to finance a transfer payment or benefit. Thus the indirect tax system becomes effectively a combined tax and transfer system, such that the same *net* amount of tax revenue is collected as when food is zero-rated, but there is a universal transfer payment

8 This is consistent with the elasticities converging towards unity, though the convergence may not be uniform. The budget shares are also affected by the tax structure, since $\dot{w}_i = \tau_i + \sum_r e_{i,r}\tau_i$, where $e_{i,r}$ is the cross-price elasticity of demand for good i with respect to a change in the price of good r.

9 On reranking arising from indirect taxes in Australia, see Creedy (2002).

of b, financed from the tax rate of τ applied to all goods. The transfer is not subject to income taxation, but of course the expenditure arising from b is taxed at the standard GST rate. This may be regarded as the least redistributive option, compared with a benefit structure which more closely targets poor goups. The government budget constraint, since T_E must be obtained net of the transfer, is thus given by:

$$\tau(\bar{y}+b)-b=T_E \tag{7}$$

and the transfer is:

$$b=\frac{\tau\bar{y}-T_E}{1-\tau}=\frac{\tau(\gamma+\alpha\bar{y})}{1-\tau} \tag{8}$$

The average tax rate in this structure, ATR_B, is thus found to be:

$$ATR_B=\frac{\tau(y+b)-b}{y}=\tau-\frac{b}{y}(1-\tau) \tag{9}$$

Substituting for b gives:

$$ATR_B=\tau\left\{1-\frac{\gamma+\alpha\bar{y}}{y}\right\} \tag{10}$$

Comparison of the average tax rates in the two alternatives gives:

$$ATR_B-ATR_E=\alpha\tau\left(1-\frac{\bar{y}}{y}\right) \tag{11}$$

and $ATR_B-ATR_E>0$ for $y>\bar{y}$, and is <0 for $y<\bar{y}$. Hence, just as with a 'basic-income flat-tax' income tax and transfer system, redistribution takes place from those above the arithmetic mean to those below the mean total expenditure. Even though the transfer is universal, and thus received by those with high and low total expenditures, it is more redistributive than a system that exempts food. This is also shown clearly by the fact that the average tax rate increases more rapidly with the transfer than with exemptions over the whole range of total expenditures, since:

$$\frac{dATR_B}{dy}=\frac{\tau(\gamma+\alpha\bar{y})}{y^2}>\frac{dATR_E}{dy}=\frac{\gamma\tau}{y^2} \tag{12}$$

for all y.

The comparisons in this section relate to progressivity measures defined in terms of expenditure, such as the average tax rate. However, it is also possible to consider measures based on welfare changes in relation to total expenditure (such as the ratio of the equivalent variation to total expenditure). Such measures allow for demand changes arising from relative price changes. It can be shown that, if there is some heterogeneity in preference parameters, welfare and net expenditure measures of the effects of tax and price changes do not necessarily agree; see Creedy and Sleeman (2006). For the empirical work presented below, demand responses are modelled and welfare changes are measured in terms of equivalent variations.

3 Demand Responses and Categories

The demand elasticities used below were obtained using a result established by Frisch (1959) for directly additive utility functions, which relates own- and cross-price elasticities to total expenditure elasticities, budget shares and the elasticity of the marginal utility of income (the so-called ‘Frisch parameter’). There are well-known criticisms of the use of additivity; for example it does not allow for complements. Where welfare changes are obtained, these are based on the Linear Expenditure System (LES), a special case of an additive utility function. In view of these strong assumptions, the results must be treated with caution. The alternative is to make the unrealistic assumption that demand patterns are fixed when prices change.

The approach involves a set of price elasticities being computed for each of a range of total expenditure groups for several household types, following the general approach suggested in Creedy (1998b), and described briefly in the appendix.

Instead of using a single set of parameters, estimates of the LES are obtained for each household type for each of a number of total expenditure groups. Households within each group are assumed to have the same preferences, but these are allowed to vary between groups.[10] Hence many elasticities are computed.

3.1 The Clustering Technique

The present paper makes use of the 12 representative household types, referred to as clusters, identified by Ball and Ryan (2014). The clusters were formed by grouping together households that were found to be similar on a range of economic and demographic dimensions. These dimensions include the age of highest income earner, number of children, qualification, home ownership status (mortgage holder or renter), and household disposable income. Additional dimensions are the proportions of income that comes from the following sources: government transfers (excluding Working for Families); private and public pensions; investments and private sources (excluding private pension and investments). The clusters were created by applying the $K-$Harmonic Means algorithm to the

[10] If all households have identical tastes, additivity implies that optimal indirect tax rates are uniform. However, this does not arise in the present context because of the allowance for heterogeneity of expenditure patterns between groups.

2006/07 Household Economic Survey dataset. The clustering algorithm, following Ball and Ryan (2014), is briefly described as follows.

For a given household, for each dimension, d, define the index, x_d, using:

$$x_d = \omega_d a_d \tag{13}$$

where ω_d is the weight assigned to the importance of dimension d and (if the values that dimension can take are numeric) a_d is the observed value of that dimension for the given household, standardised to a value between 0 and 1 based on its percentile relative to all observed values of that variable in the dataset. For $1, 2, \ldots, d$ dimensions there is a vector:

$$X = (x_1, x_2, \ldots, x_d) = (\omega_1 a_1, \omega_2 a_2, \ldots, \omega_d a_d) \tag{14}$$

that describes each household; there is also a vector M_i:

$$M_i = (m_{i1}, m_{i2}, \ldots, m_{id}) \tag{15}$$

of the values of the index for each dimension d at the centre of cluster i. Let:

$$D(\Omega, M) = \sum_{X \in \Omega} d(X, M) \tag{16}$$

be the distance measure that describes the distance between each observation X in the whole dataset Ω and all the centres M, summed across all observations. We find the cluster centres that minimise the following distance measure:

$$D(\Omega, M) = \sum_{X \in \Omega} \left(\frac{K}{\sum_{i=1}^{K} \|X - M_i\|^2} \right) \tag{17}$$

where K is the number clusters created (in our case 12). As can be seen with the inside summation over i this measure considers the distance from every observation X to the centre of each cluster M_i. [11]

3.2 Cluster Descriptions

This section provides a brief description of the 12 clusters created using the K-Harmonic Means, with Table 1 providing some key income and demographic information. Ball and Ryan (2014) described clusters A and B as renting young low-income households. Furthermore, B has the highest number of children on

[11] For a more detailed exposition of the clustering methodology used, including information on the dimension weights and the methods used to select the number of clusters, as well as a technical discussion of the appropriateness of clustering for the HES data set, see Ball and Ryan (2014).

average of any cluster (3.0) which, coupled with their relatively low income, means they have the highest food budget share of any cluster. In terms of explaining the differences in food budget shares for the two middle aged mortgage holding clusters, I and J, the average number of children appears to matter again. Cluster I has higher average wage and salary income than J, but has more children and hence a lower disposable income per adult equivalent, and a higher food budget share of 14.05 compared with 12.31 for cluster J.

Table 1: Cluster 2010 Demographic and Income Information

	Average within each cluster					
	HH disp income per AE	Wage and salary income	Population (000s)	Age of HH head	No of children	Food share
A	21909	23229	156	33	0.7	13.82
B	22977	33986	104	38	3.0	19.38
C	25838	33076	132	34	0.6	12.27
D	16200	10296	103	57	0.1	15.94
E	20938	2266	172	75	0.0	17.62
F	33168	37952	138	52	0.2	12.73
G	60640	101020	159	30	0.3	11.13
H	27455	6315	115	72	0.0	16.18
I	45461	88866	138	42	2.5	14.05
J	52247	84086	163	47	0.6	12.31
K	71284	53794	87	63	0.2	13.71
L	72624	108176	158	53	0.3	10.99

Clusters E, H and K are the older household clusters. Cluster E contains those who are fully retired and receive New Zealand Superannuation (NZS). Those in H work part-time in retirement while receiving NZS. Cluster K contains full-time workers receiving NZS or those still working but nearing the age of eligibility for NZS. Of these clusters, E has the highest food budget share (and the second highest of any cluster) reflecting their lower income. Those in K, who are working full-time and hence have the highest disposable income, have the lowest food budget share of the three older clusters. There are two young highly qualified clusters, C and G. Cluster G is characterised as being a cluster of young well-paid professionals, with relatively high salary earnings and a low average number of children; households in cluster C, although similarly qualified as G, differ from G in-so-much as they receive lower wage and salary income. Reflecting their high income and lack of children, cluster G has a relatively low budget share of food (the second lowest of any cluster), whilst cluster C's budget share is low relative to other clusters reflecting their lack of children but slightly higher than G's reflecting their lower income. The inverse relationship between income and food budget share holds in the clusters with average age 50 to 60 as well. Members of L, mid-life highly qualified high earners with high investment income, have a higher relative budget share on luxury items, such as international air travel, audio-visual and computing equipment and major cultural and recreational equipment but the lowest budget share of any cluster on food (11 per cent). Households in Cluster F, on average, are roughly the same age as L, slightly less qualified and on lower incomes, and have a food budget share of 12.73 per cent compared with 10.99 per cent. The lowest income cluster of these, D, a beneficiary cluster, with relatively large average payments of unemployment, invalid and sickness benefit, relatively low wage and salary income, and a food budget share of 15.94 per cent.

4 Welfare Effects of GST Changes

The welfare effects reported here are for two policy changes. The first change simply eliminates GST on food, so that all households are expected to gain. The second policy involves the zero-rating of food while at the same time keeping total GST revenue constant by raising the GST rate imposed on all other goods. The measure of welfare change is the equivalent variation (EV), defined as the maximum amount an individual would be prepared to pay (after the tax policy change) to return to the old prices. Hence a negative value of EV indicates a welfare gain rather than a loss. The method of calculation is described in the Appendix.

For comparisons among households with differing compositions, parametric adult equivalent scales are used such that the adult equivalent size, s, of a household with n_a adults and n_c children is given by $s = (n_a + \theta n_c)^{\alpha}$, with $\theta = 0.6$ and $\alpha = 0.8$. The latter reflects the assumed extent of economies of scale within the household. The data used are obtained from the Household Economic Survey for 2010.

The empirical results are presented in Table 2 for the clusters and in Table 3 for the deciles of the distribution of household expenditure per adult equivalent person.[12] Consider first the section of the table for deciles of household expenditure per adult equivalent person. The final column shows the average budget share for food within each decile.[13] This decreases systematically as total expenditure

[12] The unit is the household: that is, 10 per cent of households are in the lower decile of household expenditure per adult equivalent.

[13] All values reported are averages within the designated groups, obtained as weighted averages using the HES sample weights used to aggregate to population values. In the case of the budget shares, these are obtained as the ratio of the (weighted) average household expenditure on food, divided by average total expenditure. This is in line with Statistics New Zealand calculations of budget shares (though of course the ratio of averages is not the same as the average of ratios).

Table 2: Welfare Changes and Budget Shares for Clusters

	No GST on food		Revenue neutral		
	EV	EV/y	EV	EV/y	Food share
Clusters					
A	-534	-1.68	-15	-0.05	13.82
B	-1028	-2.19	-292	-0.62	19.38
C	-696	-1.53	59	0.13	12.27
D	-486	-1.88	-70	-0.27	15.94
E	-456	-2.16	-125	-0.59	17.62
F	-733	-1.54	55	0.12	12.73
G	-936	-1.32	260	0.37	11.13
H	-663	-1.92	-112	-0.32	16.18
I	-1418	-1.66	-22	-0.03	14.05
J	-1026	-1.51	105	0.15	12.31
K	-1011	-1.56	63	0.10	13.71
L	-1139	-1.36	267	0.32	10.99

increases, consistent with Engel's Law. The first policy change considered is the simple zero-rating of food. With a Goods and Services tax-exclusive rate of 0.125, or 12.5 per cent (consistent with the use of the 2009/10 Household Expenditure Survey), this implies a reduction in the price of food, while the prices of other goods remain constant in this partial equilibrium setting.[14] Hence it is expected that all households are better off as a result of this change, as indicated by the fact that all values of EV in the second column of the table are negative: the convention in the public finance literature that a welfare loss is indicated by a positive EV is followed here. Although all households gain, it is clear that those in the top decile gain considerably more than those in the lower decile in absolute terms. This reflects the poor 'target efficiency' of this kind of policy change where it is intended to generate greater equality: it is impossible to exempt food from indirect taxation for only those in specific categories defined by personal characteristics. Nevertheless, in relative terms, the lower deciles gain relatively more than the higher deciles: the ratio EV/y declines systematically as y increases, from -2.69 in the first decile to -1.04 in the top decile.[15]

This kind of tax policy change is not, however, revenue neutral. In practice it would be necessary to raise tax revenue in other ways or reduce expenditure, in order to avoid debt increases. Suppose that the extra revenue (in aggregate) is obtained

[14] It could be argued that, since food may enter as an intermediate input into some goods and services which are still subject to GST, those other goods will not have constant prices. This kind of effect cannot be considered here but can reasonably be expected to be small.

[15] There is an exception in that the average ratio for decile 4 is slightly higher than that for decile 3.

Table 3: Welfare Changes and Budget Shares for Deciles

	No GST on food		Revenue neutral		
	EV	EV/y	EV	EV/y	Food share
Equivalised expenditure deciles					
Decile 1	-472	-2.69	-206	-1.18	24.16
Decile 2	-619	-2.38	-216	-0.83	20.09
Decile 3	-620	-1.95	-108	-0.34	17.12
Decile 4	-759	-1.99	-147	-0.38	16.79
Decile 5	-796	-1.78	-65	-0.15	15.47
Decile 6	-899	-1.74	-52	-0.10	15.39
Decile 7	-966	-1.59	40	0.07	13.45
Decile 8	-998	-1.52	99	0.15	12.66
Decile 9	-1003	-1.32	292	0.38	11.19
Decile 10	-1188	-1.04	803	0.70	7.87

by increasing the GST rate applied to taxed goods and services. It is found, by a process of trial and error, that the rate would need to increase by 2 percentage points. This is the rate obtained after all households adjust their consumption pattern as a result of the relative price change between taxed and untaxed goods. The effects of this revenue neutral change are shown in the two columns under the heading 'Rev neutral' in the table. For the first six decile groups the EV is still negative, reflective a welfare gain, but for the higher-decile groups there is a welfare loss. When converting to the ratio, EV/y, the revenue neutral policy change is progressive in that the absolute value of the ratio, EV/y, decreases as y increases. However, the relative gains to the lower-decile groups are much smaller than with the first (revenue decreasing) policy change.

Turning to the results for the clusters, whose characteristics are summarised in Table 1, all groups gain from the aggregate revenue-reducing policy of simply zero-rating food. Those with the largest gains from this policy, clusters A, B, D, E, H and I, continue to gain from the revenue-neutral policy. However, as before, their relative gains are much lower. Clusters D and E are the older households with few, if any, children and low average incomes, with relatively large food budget shares. Clusters B and I are households with relatively more children and, as expected, relatively high budget shares for food. The biggest losers from the revenue-neutral policy are clusters G and L, consisting of high-income households with few children and hence lower food budget shares.

Table 4 shows results for various other groups distinguished by age, household

Table 4: Welfare Changes and Budget Shares for Various Groups

	No GST on food		Revenue neutral		
	EV	EV/y	EV	EV/y	Food share
Age					
$<$25	-720	-1.45	133	0.27	11.56
25-34	-852	-1.55	79	0.14	13.01
35-44	-977	-1.71	-21	-0.04	13.56
45-54	-1011	-1.55	101	0.15	12.77
55-64	-850	-1.50	116	0.20	12.61
65+	-622	-1.98	-109	-0.35	16.21
Household structure					
Single	-409	-1.53	43	0.16	12.07
Single with children	-693	-1.77	-46	-0.12	14.64
Couple only	-912	-1.54	89	0.15	12.66
Couple with children	-1222	-1.78	-89	-0.13	14.33
Other: no children	-1086	-1.67	2	0.00	14.63
Other: with children	-841	-1.43	162	0.28	12.25
Housing tenure					
Mortgage	-892	-1.61	45	0.08	13.12
Renting	-737	-1.66	8	0.02	13.79
Other	-1020	-1.61	47	0.08	13.13

structure, and housing tenure.[16] In particular it is seen that, with the groups classified by household structure, those with children are generally the only ones to experience welfare gains from the revenue neutral policy change. Consistent with this finding, the age group 35-44 – the prime child-rearing age bracket – experience welfare gains from the revenue neutral change. Again, not surprisingly, those in the age group 65 and over also gain, on average, from the revenue neutral policy.

These results demonstrate clearly that the absolute and relative gains and losses from a revenue neutral policy of zero-rating food in a GST are small relative to total expenditure, despite the fact that the policy can achieve some progressivity. Such exemptions are a 'poor redistributive instrument'. As demonstrated above using a simple illustration, a policy of raising transfer payments – even where these are received by everyone – is capable of producing more progressivity.

The above results may be compared briefly with those obtained by a simple comparison of tax payments in relation to total expenditure, with no allowance

[16] Values for 'Other households' are omitted from the table. These two groups (with and without children) are extremely heterogeneous, and in the case of those without children the EV is negligible, making the ratio of the change in tax to the equivalent variation extremely high.

for demand responses to price changes, instead of using equivalent variations. Table 5 shows the ratio of the reduction in GST to the equivalent variation, for the revenue neutral policy change which, as explained above, involves an increase in GST of 2 percentage points. Where this ratio has a positive sign, the tax changes and welfare changes are in the same direction. A negative sign indicates that the use of tax changes, with no demand responses, suggests a change in the opposite direction from that of the welfare change. In all the cases where there is a negative sign, the use of tax changes indicates a gain, whereas the use of equivalent variations indicates a welfare loss. Even where the two measures agree about the direction of change, it is clear that there are substantial differences in relative orders of magnitude.

Table 5: Ratio of Tax Change to Equivalent Variation: Revenue Neutral

Age	ΔTax/EV
<25	0.85
25-34	0.28
35-44	0.85
45-54	0.45
55-64	0.44
65+	1.07
Household structure	
Single	0.99
Single with Children	1.44
Couple Only	0.57
Couple with Children	0.99
Other: no children	-45.50
Other: with children	0.50
Tenure	
Mortgage	0.32
Renting	-3.15
Other	0.32

Clusters	ΔTax/EV
A	1.34
B	1.25
C	1.05
D	1.25
E	0.93
F	0.65
G	0.77
H	1.13
I	3.67
J	0.84
K	-0.52
L	0.94
Deciles	
Decile 1	1.18
Decile 2	1.05
Decile 3	1.43
Decile 4	1.16
Decile 5	1.89
Decile 6	2.63
Decile 7	-0.25
Decile 8	0.56
Decile 9	0.71
Decile 10	1.00

5 Conclusions

This paper has investigated the detailed welfare effects on households of a change to the indirect tax structure in New Zealand involving the zero-rating of food in the GST. Such a policy is typically motivated by a desire to introduce some progressivity into the GST structure, in view of the well-established property that the share of household expenditure devoted to food declines as household income increases. As mentioned in the introduction, it is sometimes argued that the case for zero-rating is greater, the higher is the uniform GST rate. However, richer households benefit more in absolute terms. The poor targeting of such a policy is illustrated in a simple model where the revenue lost from zero-rating food is instead devoted to a universal transfer payment, with a larger effect on progressivity.

The detailed effects of zero-rating food, for a range of household types, were then investigated using Household Economic Survey data. Demand responses to consumer price changes were estimated and welfare changes, in terms of equivalent variations, were obtained. Comparisons were made across ‘clusters’, consisting of groups of households with similar characteristics, as well as for households in decile and age groups. A revenue neutral reform is seen to produce a very small amount of progressivity in the GST, with welfare losses for higher decile groups and welfare gains for the lower deciles. Furthermore, redistribution is from richer households without children to poorer households with children, and older households. It was also shown that the use of tax changes, with an assumption of no responses to price changes, somewhat overstates the degree of progressivity achieved. The analysis supports earlier studies suggesting that indirect tax exemptions and zero-rating provide a poor redistributive instrument compared with direct taxes and transfers.

Appendix A: Utility, Demand Elasticities and Welfare

This appendix describes the method used to obtain demand elasticities and welfare changes using the linear expenditure system (LES), applied separately for a range of demographic groups (though the following notation generally omits the additional subscript).

The first stage is to obtain, for each household type, a set of average budget shares, w_{ki}, for each consumption category, i, and a range of total expenditure groups, k. The total expenditure elasticities are obtained using the variations in budget shares for each commodity group. However, the observed variability in budget shares gives rise to some negative total expenditure elasticities. This can be overcome by smoothing the data. The approach used was first to carry out a series of ordinary least squares regressions of the form:

$$w_{ki} = a_i + b_i \log y_k + c_i \left(1/y_k\right) \tag{A.1}$$

for each commodity group (and household type), where the values of y_k correspond to the arithmetic mean values of total expenditure in each group, k. The form in (A.1) provides a reasonably good fit for most groups and ensures that the predicted weights add to unity.

The second stage is to compute own- and cross-price elasticities, e_{ii} and e_{ij}, (again for each total expenditure group and household type) using Frisch's (1959) results for additive demand systems. The expressions require the use of the elasticity of the marginal utility of total expenditure with respect to total expenditure, ξ, often referred to as the 'Frisch parameter'. If δ_{ij} denotes the Kronecker delta, such that $\delta_{ij} = 0$ when $i \neq j$, and $\delta_{ij} = 1$ when $i = j$, Frisch showed that the elasticities can be written as:

$$e_{ij} = -e_i w_j \left(1 + \frac{e_j}{\xi}\right) + \frac{e_i \delta_{ij}}{\xi} \tag{A.2}$$

It is necessary to make use of extraneous information about the way in which the Frisch parameter varies with total expenditure. In view of the role played by the Frisch parameter and the lack of a really firm foundation for the values used, it is important to carry out a range of sensitivity analyses.

The third stage involves obtaining parameters of the LES direct utility function (again for each total expenditure group and household type):

$$U = \prod_i (x_i - \gamma_i)^{\beta_i} \tag{A.3}$$

with $0 \leq \beta_i \leq 1$; γ_i is the committed consumption of good i, and $\sum \beta_i = 1$. The own-price elasticity, e_{ii}, is given by:

$$e_{ii} = \frac{\gamma_i (1 - \beta_i)}{x_i} - 1 \tag{A.4}$$

The total expenditure elasticity of good i, e_i, is:

$$e_i = \frac{\beta_i y}{p_i x_i} \tag{A.5}$$

Having obtained the total expenditure elasticities from the smoothed budget shares, the corresponding values of β_i at each total expenditure level were obtained using (A.5), whereby $\beta_i = e_i w_i$. Using the values of own-price elasticities as described in the second stage above, equation (A.4) can be used to solve for $p_i \gamma_i$, the committed expenditures for each good.

As before, the various parameters vary with y, but the additional subscript is suppressed for convenience. Defining the terms A and B respectively as $\sum_i p_i \gamma_i$ and $\prod (p_i/\beta_i)^{\beta_i}$, the indirect utility function for the LES, $V(p, y)$, is:

$$V = (y - A) / B \tag{A.6}$$

The expenditure function, $E(p, U)$, the minimum expenditure required to achieve U at prices p, is found by inverting (A.6) and substituting E for y to get:

$$E(p, U) = A + BU \tag{A.7}$$

If the vector of prices changes from p_0 to p_1, the equivalent variation, EV, is $EV = E(p_1, U_1) - E(p_0, U_1)$. Substituting for E using (A.7) gives:

$$EV = y - (A_0 + B_0 U_1) \tag{A.8}$$

Substituting for U_1, using equation (A.6) into (A.8) and rearranging gives:

$$EV = y - A_0 \left[1 + \frac{B_0}{B_1} \left(\frac{y}{A_0} - \frac{A_1}{A_0} \right) \right] \tag{A.9}$$

The term A_1/A_0 is a Laspeyres type of price index, using γ_is as weights. The term B_1/B_0 simplifies to $\prod (p_{1i}/p_{0i})^{\beta_i}$, which is a weighted geometric mean of price relatives. These two terms can be expressed in terms of the $\dot{p}$s.[17] If all prices change by the same proportion, $\dot{p}_i = \dot{p}$ for all i, and (A.9) becomes $EV/y = (1 - B_0/B_1) + (A_0/y) \{(B_0/B_1)(A_1/A_0) - 1\}$, with $B_1/B_0 = A_1/A_0 = 1 + \dot{p}$.

[17] Since $p_{1i} = p_{0i}(1 + \dot{p}_i)$, and defining $s_i = p_{0i}\gamma_i / \sum_i p_{0i}\gamma_i$, it can be shown that $A_1/A_0 = 1 + \sum_i s_i \dot{p}_i$ and $B_1/B_0 = \prod_i (1 + \dot{p}_i)^{\beta_i}$.

References

Ball, C. and M. Ryan (2014). New Zealand households and 2008/09 recession. *New Zealand Economic Papers 48*(1).

Creedy, J. (1998a). Are consumption taxes regressive? *Australian Economic Review 31*(2), 107–116.

Creedy, J. (1998b, June). Measuring the Welfare Effects of Price Changes: A Convenient Parametric Approach. *Australian Economic Papers 37*(2), 137–51.

Creedy, J. (1999). Marginal indirect tax reform in australia. *Economic Analysis and Policy 29*, 1–14.

Creedy, J. (2001). Indirect tax reform and the role of exemptions. *Fiscal Studies 22*(4), 457–486.

Creedy, J. (2002). The GST and Vertical, Horizontal and Reranking Effects of Indirect Taxation in Australia. *Australian Economic Review 35*(4), 380–390.

Creedy, J. and N. Gemmell (2004). The built-in flexibility of income and consumption taxes in New Zealand. *Australian Economic Papers 43*(4), 459–474.

Creedy, J. and C. Sleeman (2006). *The Distributional Effects of Indirect Taxes: Models and Applications from New Zealand*. Edward Elgar.

Frisch, R. (1959). A complete system for computing all direct and cross demand elasticities in a model with many sectors. *Econometrica 27*, 177–196.

Keen, M. (2013). The anatomy of the VAT. *National Tax Journal 66*, 423–446.

Stern, N. (1990). Uniformity versus selectivity in indirect taxation. *Economics and Politics 2*(1), 83–108.

van Klink, A. (2012). How to zero-rate the GST on food: best and worst practice from the United Kingdom, Canada and Australia. *New Zealand Journal of Taxation Law and Policy 18*, 276–302.