

Does Demand Create Poor Quality Supply: A Critique of Alternative Distributional Analyses*

This paper was prepared by Dr Michael Carnahan, Taxation Policy Group, Commonwealth Treasury for presentation at the Economic Modelling Bureau of Australia conference on Tax Reform on 1 September 1998. The paper was not presented as the calling of the Commonwealth election placed the Government in caretaker mode, precluding public comment by officials. The paper was released publicly by the Treasurer on 8 October 1998.

**The author acknowledges the contributions of Ken Henry, David Tune, Alan Preston, Steve Hatfield Dodds, Harry Kroon, Keith Woolforde and Treasury and ABS participants at a Taxation Policy Group seminar for helpful comments — the usual caveats apply. Views expressed in this paper do not necessarily represent the views of the Commonwealth Government or the Treasurer.*

INTRODUCTION

The release of any major reform proposal creates considerable uncertainty amongst the community. Economic commentators and the public more generally have two distinct but related questions in their minds.

- Will I gain or lose from this reform (in both a relative and absolute sense)?
- In the community more widely, how many people will win or lose, and are they in groups that I care about?

In the vast majority of cases, and for the vast majority of reform proposals, these questions cannot be answered definitively. However, what can be predicted with certainty is that there will be a rekindling of interest in distributional analysis. When such reforms are announced, the demand for distributional analysis and commentary skyrockets. There is an almost insatiable demand for detailed, disaggregated analysis, particularly by the media and by opponents of reform. The former are interested in results that can be translated into banner headlines; the latter seek support for the view that they or their constituents will be harmed by the proposals.

Much of the analysis generated in response to this demand comes heavily caveated, and where it doesn't, it should. While these shortcomings and caveats may render much analysis valueless, the caveats are rarely, if ever, reported.

Instead, the caveats are relegated to technical appendices (readers interested in data shortcomings are referred to other sources) and the results are reported without caveat and given an imprimatur of accuracy they generally do not deserve.

The key question is one that is familiar to any student of estimation theory: what approach gives the most useful information about the issue of interest? And like most issues in estimation theory, the problems turn on data shortcomings and determining what can be usefully and defensibly concluded from generally poor data.

The major conclusion in this paper is that the **best** estimate of the impact on a household of a change in prices is the population CPI impact, and the **best** estimate of household saving rates is either the population saving rate or zero. Hence, this approach was adopted in the distributional analysis undertaken in the Government's tax reform document *A New Tax System*. The purpose of this paper is to provide a critique of this approach and the alternatives.

Whether these are **good** estimates is a related, but separate question. In the absence of useable disaggregated data on expenditure patterns, these are the best estimates. Underpinning this is the fact that the use of the Household Expenditure Survey (HES) for disaggregated distributional analysis is invalid. This is in some ways not surprising, as the survey is not designed for such a task.

The remainder of the paper is divided into three sections. The next section outlines what measure should be used, at a conceptual level, to determine whether somebody, or a group wins or loses from tax reform. The net cash gain measure presented in this section is used frequently by those undertaking distributional analysis. In the following two sections the two major contested issues are addressed: how should the impact of price changes be estimated; and how should saving rates be estimated.

WHAT DEFINES A WINNER?

Whenever major changes take place there is a demand for a single, easily understood figure — generally a dollar amount — to be produced to fully capture the impact of the reform on any particular household. There are a number of possible measures that can be presented to estimate the impact of a tax reform package. Major tax reform proposals have an impact on both income (from labour and capital) and prices. The conceptually pure measure presented in the economic textbooks is the compensating variation. A more tractable measure is the net cash gain figure that has been used in the analysis of previous reform proposals. Within the net cash gain framework there are questions relating to measuring price impacts and what saving rate to apply.

Compensating Variation

Compensating variation¹ measures the amount of money that can be taken from a household after a change in tax inclusive prices (p) and disposable income (Y_d), leaving the household as well off as it was before the changes. Formally the compensating variation (CV) is defined implicitly in the following equation, where v is the household's indirect utility function and the superscripts 0 and 1 denote pre and post reform.

$$v(\underline{p}^0, Y_d^0) = v(\underline{p}^1, Y_d^1 - CV) \quad (1)$$

This measure obscures an interesting issue. Disposable income is treated as a parameter, rather than a choice variable in this framework. However, a household's disposable income depends on the labour supply and investment decisions of the household². So it would be more correct to make v a function of the after tax wage rate, after tax return on capital and output prices.

The major insight gained from the concept of compensating variation is that the net cash gain measure discussed below understates the gains (or overstates the losses) associated with reform. When prices change, people alter their decisions — buying more of the goods and services whose prices have dropped and less of the items that have become more expensive. So if prices change and the household receives an amount of compensation that just allows them to purchase the bundle of goods and services they previously purchased, they will be overcompensated. This issue comes even more clearly into focus when labour supply decisions are included in the analysis.

Net Cash Gain Measures³

Problems in estimating the indirect utility function generally mitigate against the use of compensating variation as a measure of the effect of reform. Instead, a range of measures have been developed to estimate how much extra cash households have after any changes. The simplest of these derives from the compensating variation measure above and calculates the change in disposable income and the change in the cost of buying the goods and services purchased before the reform. Using the notation as above, denoting the net cash gain by

1 This discussion is intentionally cursory and summary. A more detailed discussion can be found in any advanced microeconomics text (eg Deaton and Muellbauer (1980)).

2 The disposable income will also depend on the initial asset endowment of the household, but these can be more accurately represented as parameters, particularly if intertemporal choices are included.

3 This discussion intentionally skirts around several issues including whether gains should be measured in pre or post reform dollars and what is happening to the stock of savings. These issues are not unimportant. However, the focus of the paper is on estimation and any measure of gain allows the estimation issues discussed below to be teased out.

NCG and letting x be the vector of goods and services purchased by the household, the measure is given by:

$$NCG = (Y_d^1 - Y_d^0) - (\underline{p}^1 \cdot \underline{x}^0 - \underline{p}^0 \cdot \underline{x}^0) \quad (2)$$

Equation (2) can be rewritten as:

$$NCG = (Y_d^1 - Y_d^0) - \left(\frac{\underline{p}^1 \cdot \underline{x}^0}{\underline{p}^0 \cdot \underline{x}^0} - 1 \right) \cdot \underline{p}^0 \cdot \underline{x}^0 \quad (3)$$

Noting that disposable income is either spent or saved, and defining s as the proportion of disposable income that is saved, it is possible to rewrite equation (3) as:

$$NCG = (Y_d^1 - Y_d^0) - \left(\frac{\underline{p}^1 \cdot \underline{x}^0}{\underline{p}^0 \cdot \underline{x}^0} - 1 \right) \cdot (1 - s) \cdot Y_d^0 \quad (4)$$

But $(\underline{p}^1 \cdot \underline{x}^0 / \underline{p}^0 \cdot \underline{x}^0)$ is a standard Laspeyres price index measure⁴. The official CPI calculated by the ABS uses the Laspeyres price index methodology⁵. Hence the net cash gain can be rewritten as:

$$NCG = (Y_d^1 - Y_d^0) - CPI \cdot (1 - s) \cdot Y_d^0 \quad (5)$$

At a conceptual level, the material presented to this point is reasonably non-controversial and not particularly revelatory. Household disposable income is a scalar variable and can be calculated simply and definitively if private income and benefit entitlements are known (and assuming entitlements are accurately accessed). However obtaining information on the actual level of household income is difficult. Moreover, the impact of the relative (and absolute) price changes, ie the estimation of the CPI, and the estimation of the saving rate are considerably more contested. These estimation issues are discussed in the following two sections.

ESTIMATING PRICE IMPACTS

The central issue in estimating price effects is determining what approach or measure provides the most useful information to somebody trying to interpret the impact of a policy package on different households. That is, like any

⁴ For a more detailed discussion of price indexes and cost of living measures see, for example, Deaton and Muellbauer (1980) Chapter 7.

⁵ For a precise definition of the official CPI see Australian Bureau of Statistics (1993).

estimation problem, the choice is constrained, and the aim is to get the most informative measure — in the absence of the correct measure.

Three approaches have been used to capture the effect of relative (and absolute) price changes.

- Calculate a single price impact for all households using the aggregate consumption patterns contained in the CPI to capture the effect of relative price changes.
- Calculate different price impacts for different household groups.
- Calculate individual price effects for each record in a sample of the population, and then estimate the impact on the entire population based on this sample.

The first two approaches lend themselves to the production of net cash gain charts or tables for different household groups. These approaches are particularly useful for generating sensationalist headlines along the lines of ‘Group X lose \$20 p/w’. They also allow all readers to think that they can determine how much they themselves will gain by reading the number off the chart. The third approach lends itself to the direct calculation of winners and losers, generating headlines such as ‘200 000 battlers lose’.

There are arguments for adopting the most highly disaggregated approach, as different people do indeed buy different things. Conversely, if the time period over which the spending patterns are observed is too short, there may be arguments for adopting a more aggregated approach. Over a week, two households may have very different spending patterns, but over a month or a year or ten years their spending patterns may converge.

While interesting, the conceptual issues are swamped by the empirical issues. Specifically, data limitations invalidate any individual based analysis, ruling out the third estimation approach (household specific effects). These limitations also cast very serious doubts on the validity of the second approach (analysis by household group) — particularly if the group size is small. The data problems are discussed below, followed by an assessment of each of the three approaches in the light of these limitations.

The Use of Household Expenditure Survey Data in Distributional Analysis⁶

The Household Expenditure Survey conducted by the Australian Bureau of Statistics has frequently been the source of the disaggregated data used to calculate household group-specific or household-specific CPI effects. However, the fact that the HES is a sample survey, and other features of the sample design,

⁶ This section draws on Wright and Dolan (1992).

significantly constrain the scope for such disaggregation. By its very nature, the division of HES data into the smaller subsets required for distributional analysis rapidly erodes the reliability of estimates thus derived. Related to this, the HES sample has not been designed to permit the calculation of reliable representative results for all of the household types for which price effects might be desired.

Treatment of durable good purchases

A separate constraint on disaggregation arises from the treatment of expenditure on durables. The HES records expenditure on home purchases, motor vehicles and other consumer durables as it occurs. As a corollary to this treatment, the HES records sales of houses and consumer durables as negative expenditure. The sample is so designed that the inclusion of lumpy expenditures incurred in a relatively short reporting period provides a reasonable proxy for the cost of the ongoing service flows associated with those durable expenditures — **provided one is concerned with the average expenditure of the full sample**. For example, the purchase of a house by one household is balanced by sales by other households and the lack of such expenditure for many other households during the relatively short reporting period of the HES. When the sample average is calculated, the average expenditure on housing services is believed to be representative on an ongoing or weekly basis. A simple numerical example is presented in Box 1.

The acquisitions approach to the reporting of expenditures on durables necessarily implies that household-specific estimates of expenditure are intrinsically inaccurate. The implication is that the larger the degree of disaggregation required by distributional analysis, the less reliable will expenditure estimates for durables be (irrespective of the sampling errors otherwise introduced by sample subsets). In effect, HES is designed to provide reliable estimates of household expenditure on durables when averaged across the entire sample but does so in a manner that precludes reliable household-specific estimates and casts considerable doubt on the reliability of group-specific estimates.

Inconsistency between HES and official CPI spending patterns

In addition, there is a mutual inconsistency between the expenditure patterns implied by the HES and those implied by both the National Accounts and the official ABS CPI series. The extent of some of these inconsistencies is illustrated in Table 1.

If the shares of some goods are below the CPI share, the shares of all other goods will be commensurately higher. Traditional understatement of the 'sin' goods (alcohol and tobacco) not only calls into question the accuracy of the HES data, as it illustrates response bias in the survey, it also corrupts other expenditure shares — where the **level** of expenditure may have been accurately reported.

Box 1: Problems with disaggregating – a simple numerical example

The approach of averaging durable good purchases across households is conceptually defensible. However, for the approach to have validity, it requires a large group size. The following simple stylised example illustrates the shortcomings arising from disaggregating the HES data. These shortcomings are greater — and are more likely to occur — the greater is the degree of disaggregation. The shortcomings associated with sample or survey bias, which manifest themselves in the inconsistency between HES and official CPI expenditure weights, are ignored in this case.

Imagine there are 52 households each receiving \$200 per week. Each week each household spends \$100 on food and saves \$100. Once a year each household spends its \$5200 saving on household furnishing and equipment. Now let the price of food rise by 10 per cent, while the price of consumer durables falls by 10 per cent. In the two week survey period ⁷, two households will purchase the durable goods, while the others will not. Any reliable distributional analysis would find a CPI effect of zero for the population and for any group (even households) within the population. However, even relatively standard distributional methodologies could produce bizarre results. Using the standard CPI/Laspeyres price index the following results are generated.

1. Single population CPI

For the population there is **no change** in the CPI.

2. Two different CPIs for two different groups

Depending upon essentially arbitrary decision rules applied by the modeller, the two households who purchased the durable good may be in the same group or in different groups.

- If one durable good purchasing household is in each group, then each group has **no change** in their CPI.
- If both durable good purchasers are in the same group, then that group has a CPI decrease of -3.3 per cent. The other group has a CPI increase of 10 per cent.

3. Different CPIs for each household

- The two households purchasing the durable good in the survey period will record a CPI impact of -9.3 per cent.
- The 50 households not purchasing the durable good in the survey period will each record a CPI of 10 per cent.

⁷ The recall period in the HES for consumer durables is actually three months, rather than two weeks. Hence the example above does not exactly replicate the HES situation. However, and most importantly, whenever the survey period does not exactly match the durable good expenditure cycle the problems discussed in the text arise.

Table 1: Comparison of CPI and HES expenditure shares⁸

Item	Percentage by which CPI share exceeds HES share	Item	Percentage by which CPI share exceeds HES share
beer/wine/spirits	91	boys clothing	-35
tobacco	59	girls clothing	66
confectionary	62	beef and veal	37
take away foods	-20	poultry	-15

The inconsistencies between the aggregate raw HES expenditure weights and the CPI series imply that the disaggregated household group data are also inconsistent.

Evaluating the three estimation approaches

Common price impact across households

Using a common price impact across households is the approach adopted in the distributional analysis presented in the Government's tax reform document *A New Tax System*. This approach is the most straightforward and easily understood. The aggregate consumption patterns represented in the official CPI are used to produce an estimate of the impact of price changes on households. The official CPI weights are based on HES expenditure shares, but the HES expenditure shares are only one input into the calculation of the CPI. While this approach may not give the sensational results demanded by some, it has several key advantages.

The CPI is the measure specifically designed by the ABS to provide an estimate of the cost of living. The methodology used to calculate the CPI figure has been refined over a number of years and is well documented and well understood. Moreover, the methodology is consistent and reasonably transparent. For these reasons the official CPI has been used historically when determining compensation for price changes. For example, increases to Government pensions and benefits have generally been determined by increases in the official CPI.

To the extent that, over time, households' expenditure patterns converge, the population CPI provides a very good estimate, at a point in time, of the impact

⁸ Drawn from Wright and Dolan (1992), p9. This analysis confines the scope of the HES to households of the type which corresponds to the CPI. Even larger differences emerge if different household types are selected from the HES.

on any individual household. Conversely, the shortcoming with this approach is that differences in consumption patterns across different groups or different households over time is not reflected. Given the considerable limitations outlined in the next two sections, the population CPI remains the best estimate of the impact on any individual household.

Different price impacts for different household groups

The second approach involves estimating different price effects for different households. There is a perception in the community that different groups have very different expenditure patterns. If this is the case, the price impact may be different for different groups, while remaining similar for members within each group⁹.

While the ABS has undertaken research and preliminary work on this topic, they do not, at this stage, have a set of CPI weights for groups within the population. This reflects both competing priorities and conceptual difficulties. Two of the conceptual problems are dealing with durable goods and dealing with under-reporting.

As discussed above, the acquisitions approach to the reporting of expenditures on durables is designed to provide reliable estimates of household expenditure when averaged across the entire sample. The extent to which this accuracy is compromised when disaggregating is an empirical issue requiring additional work. Specifically, there needs to be consideration given to what minimum sample size is needed to allow the making of meaningful inference about durable good purchases. As the CPI weights relate to expenditure **shares**, a problem with one category of expenditure will corrupt all expenditure **shares** — even though the expenditure **levels** in other categories may be well measured.

Problems associated with under-reporting were also outlined above. At a population level, under-reporting can be addressed through reference to population aggregate data from other sources — eg excise data in the case of tobacco and alcohol. This data is used in conjunction with the HES data to estimate the population CPI. However, population aggregate data — by its nature — does not source spending by sub-population group. A working assumption that under-reporting is uniform across groups is far from satisfactory given the amount of under-reporting.

One approach that has been undertaken in the past is to use the HES data as the best estimate of household spending patterns, and construct a simple Laspeyres price index based on the data contained in the HES. But such an approach neglects the shortcomings and survey and sample design issues of the HES, leading to invalid results.

⁹ A common mistake is to assume that very different spending patterns must generate very different price impacts. On the contrary, depending on what prices change and by how much, the price impacts can be very similar.

When the ABS calculate the population CPI they do not take the population HES and construct a Laspeyres price index on that data. This is because such an approach does not provide a good estimate of the change in the cost of living. Rather, in calculating the weights in the CPI series, the ABS makes a number of adjustments and normalisations, taking into account the shortcomings of the HES data with which they are dealing. Through this they construct a measure designed to estimate changes in the cost of living.

Constructing sub-population CPIs in order to measure any differences in changes to the cost of living across different groups may be possible. However, before they can be used usefully in any debate about reform packages they must be constructed as carefully as the population CPI. Moreover, the problems associated with disaggregation, including durable goods and under-reporting must also be resolved — and at this stage they are not.

Different effects for different individuals

The third approach to determining net cash gain makes more direct use of the HES. Rather than calculating a single or group CPI effect and then applying this to different income levels, the net cash gain for each of the 8 400 households is calculated using equation (2). Once this result is determined for all 8 400 households, the numbers are aggregated up by their population weights to provide an estimate of the impact of reform on the entire population.

Shortcomings with disaggregated data

This approach is superficially attractive and appears to provide the most conceptually satisfying estimate of the net cash gain for each household. However such an approach has major shortcomings which invalidate the results. The shortcomings in the use of HES expenditure data disaggregated to the individual level are discussed above. In and of themselves, these shortcomings invalidate the analysis.

As the example in Box 1 illustrates, not only is it possible to get results that seem nonsensical for some households, but the results for all other households are corrupted — not just the obviously nonsensical ones. That is, while the two households had very low CPI figures — a long way from the median — all other CPI figures were also incorrect. In addition, the median and mean were also incorrect. If **any** of the CPI figures is considered incorrect, then, as the example shows, **all** of the figures must be discounted.

What CPI results can the HES give?

Cursory analysis of HES data supports the hypothesis that they are not well designed for this disaggregated analysis. The CPI effect of the Government's tax reform package for individual HES households ranges from a high of over 100 per cent to a low of under 200 per cent. Results of this nature call into question the validity of the methodology. Arguments that results of this nature

come about from errors by those completing surveys or clerical errors hold little water.

- It is possible to come up with perfectly standard spending patterns that generate these results over the sample period but generate a CPI effect very close to a population CPI effect over a longer period.
- There is no *a priori* reason to believe that these results are any more or less valid than the results for any other household in the sample.

Correcting for ‘problem’ CPI results

As discussed above, individual household-based analysis is superficially attractive. Hence, a number of approaches have been used to ‘correct’ for the problem of outlying or ‘unrealistic’ CPI results. The two main approaches are:

- exclude households with CPI estimates above or below a tolerance level; and
- remove negative expenditures from the HES calculations.

As discussed above, and outlined in Box 1, excluding households that fall outside of a tolerance level raises more problems than it addresses. There is no more reason to remove these ‘problem’ households than there is for removing any household in the sample. Moreover, removing any household corrupts the remaining data. It is the entirety of the dataset that allows for the price impact of changes to durable good expenditures to be captured. If the households with outlying CPI effects are those buying or selling durable goods, then removing these households will remove durable goods from the expenditure bundle — leading to an upward bias on all other expenditure categories.

Moreover, the mean and median of the individual CPI figures provide little accurate information about the population CPI impact¹⁰. Hence, these data provide no indication of what would be a reasonable tolerance level. In the stylised example provided in Box 1, the two households with durable goods purchases actually had a CPI impact closer to the true CPI figure (-9.3 per cent, rather than 10 per cent) than the remaining 50 households. The two households with a -9.3 per cent CPI impact would be eliminated as outliers. The population CPI would then be found to be 10 per cent, when the correct figure is zero.

Removing negative expenditures makes as much sense as removing positive expenditures — and a HES with neither negative nor positive expenditures adds little! Both approaches seek (unsuccessfully) to address the symptoms of a flawed methodology by exclusion, rather than addressing the problems in the methodology *per se*.

¹⁰ A common mistake in this context is made by forgetting that the ratio of the means does not equal the mean of the ratios. This is clearly borne out in the stylised example presented in Box 1.

Conclusion

Returning to the question posed at the start of this section, it is clear which approach to the estimation of the impact of price changes on households is most informative. The problems associated with the HES data rule out the use of individual, household specific measures of price impacts and call into question the merit of disaggregating the 8 400 unit records into household groups. Research on the extent to which different household group spending patterns are captured by household group price indexes suggests that such price indexes add no additional value to the analysis of price impacts. Hence, analysis of the impact of price changes on households is best undertaken using a population estimate of the CPI.

ESTIMATING SAVING RATES

In principle, saving rates can also be estimated using each of the three approaches discussed above. As with price effects, saving rates will differ across households, and the question remains, what measure provides the most useful information to allow an informed assessment of the impact of a reform package on households? The shortcomings associated with the use of the HES data discussed in the previous section make the calculation of household specific saving rates problematic.

In the previous section a simple stylised example was presented to illustrate the shortcomings of calculating CPI effects from disaggregated data. This approach is extended to saving rates in this section. The results are presented in Box 2 below. Disaggregating the data can, again, provide extremely misleading results, even when the sample errors discussed elsewhere are avoided.

Shortcomings with HES income data

However, possibly even more problematic, are the shortcomings associated with using the income data collected under the HES in calculating saving rates. The ABS collects income data primarily to classify households into groups for expenditure analysis. The HES does not purport, nor attempt, to collect comprehensive income data. Hence, it is hardly surprising that the population estimates for income derived from the HES differ significantly from the ATO and DSS estimates. Of particular concern is that the estimates of tax paid and government transfer payments received generated from the HES data do not align with ATO and DSS figuring. In addition, aggregates implied by the HES, such as for wages and salaries, differ markedly from National Accounts aggregates. In the context of these shortcomings, it is not surprising that the ABS concludes:

‘HES income and expenditure estimates therefore do not balance for individual households or for groups of households and the difference

between income and expenditure can not be considered to be a measure of savings'¹¹.

Box 2: Problems with saving rates – a simple numerical example

Extending the framework outlined in Box 1, it is possible to calculate saving rates using the three distinct approaches. In the example, each household puts away half of their \$200 p/w in order to buy durable goods to the value of \$5200 once a year. In other words, each year the saving rate for all households is zero. The saving rates for the population, for groups within the population, and for the individuals in the sample are calculated using the same two week recall period as discussed in Box 1.

1. Population saving rate

For the population the saving rate is zero.

2. Two different saving rates for two different groups

The two households who purchased the durable good may be in the same group or in different groups.

- If one durable good purchasing household is in each group, then each group has a saving rate of zero.
- If both durable good purchasers are in the same group, then that group has a saving rate of -50 per cent (ie the group is dissaving). The other group has a saving rate of 50 per cent.

3. Different saving rates for different households

- The two households purchasing the durable good in the survey period will record a saving rate of -1250 per cent.
- The 50 households not purchasing the durable good in the survey period will each record a saving rate of 50 per cent.

Alternative Approaches

Given the use of the HES in calculating saving rates is not valid, there are two alternative approaches.

- Set the saving rate to zero.
- Use the population-wide saving rate implied by National Accounts data.

Setting the saving rate to zero is consistent with the approach of many researchers who take the view that current consumption (in a particular week,

¹¹ Australian Bureau of Statistics (1995), p10.

month, or even year) provides a less reliable measure than disposable income of a household's ability to consume. And it is a household's ability to consume that is really of interest, since that is what is affected by cost of living changes.

The population saving rate implied by National Accounts data is around 5 per cent. Recalling equation (4) and using this population wide saving rate, would lead to slightly higher net cash gains — assuming the price changes lead to an increase in the CPI. However, equation (4) also provides the second reason for setting the saving rates at zero. If the overall CPI effect is low, in the current case 1.9 per cent, the decision to set the saving rate at zero is not critical. If the saving rate is 5 per cent, the impact of the price increases is 1.8 per cent, rather than 1.9 per cent. Conversely, if households are dissaving at a rate of 5 per cent, the impact of price changes will be 2.0 per cent.

CONCLUSION

The demand for detailed, disaggregated analysis of the distributional impact of tax reform proposals will lead to a supply of such analysis. The analysis based on disaggregated HES expenditure patterns should be rejected, as the data cannot be validly used for this purpose. That said, there is considerable scope for additional research leading to the gathering of better data, and data that satisfies adding up tests. Such research would make a far greater contribution to the understanding of the distributional impact of current and future reform proposals than analysis based on the misuse of HES data.

BIBLIOGRAPHY

A New Tax System, (1998), Commonwealth of Australia.

Australian Bureau of Statistics (1993), *A Guide to the Consumer Price Index*, Catalogue No. 6440.0, Commonwealth of Australia.

Australian Bureau of Statistics (1995), *Household Expenditure Survey 1993-94 — User Guide*, Catalogue No.6527.0, December 1995, Commonwealth of Australia.

Deaton and Muellbauer (1980), *Economics and Consumer Behavior*, Cambridge.

Wright and Dolan (1992), '*The Use of HES data in distributional analysis*', paper presented to the Conference of Economists, July 1992.