

MELBOURNE INSTITUTE
Applied Economic & Social Research

Working Paper Series

The Welfare Implications of Unobserved Heterogeneity

Sarantis Tsiaplias

Working Paper No. 21/17
September 2017



THE UNIVERSITY OF
MELBOURNE



MELBOURNE
INSTITUTE

The Welfare Implications of Unobserved Heterogeneity*

Sarantis Tsiaplias

**Melbourne Institute: Applied Economic & Social Research
The University of Melbourne**

**Melbourne Institute Working Paper No. 21/17
September 2017**

* This study uses confidentialised unit record file data from the Household, Income and Labour Dynamics in Australia (HILDA) Survey. The HILDA Survey Project was initiated and funded by the Australian Government Department of Social Services (DSS) and is managed by the Melbourne Institute: Applied Economic & Social Research. The author thanks members of the HILDA team and participants at the 2015 Southern Economic Association (New Orleans) conference for helpful comments and suggestions. For correspondence, email <sarantis.tsiaplias@unimelb.edu.au>.

**Melbourne Institute: Applied Economic & Social Research
The University of Melbourne
Victoria 3010 Australia
Telephone +61 3 8344 2100
Fax +61 3 8344 2111
Email melb-inst@unimelb.edu.au
WWW Address melbourneinstitute.unimelb.edu.au**

Abstract

Conditions are derived for relating household well-being functions to household utility. In particular, an isomorphic relationship between the equivalent incomes stemming from subsistence-based utility functions and well-being functions is established. This allows estimates from standard models of well-being based on a CDF (eg. probit and logit models) to be given a formal welfare interpretation. New measures of the welfare distortion due to unobserved heterogeneity are also derived. An Australian household-level dataset is used as a case study for exploring the proposed measures of distortion. The results indicate that the failure to account for unobserved heterogeneity produces significant welfare distortions (primarily in the form of under-compensation). A unique welfare sensitivity curve is also estimated that indicates the presence of non-linearities that impair the typically monotonic relationship between household income, the household's capacity to adjust its income and its marginal utility of consumption. The results are significant for better understanding the welfare implications of tax and transfer policies.

JEL classification: D12, D10, C33, C35, D60

Keywords: Household expenditure, social welfare, consumption, heterogeneity, equivalence scales, income

1 Introduction

The importance of unobserved heterogeneity for welfare purposes is highlighted in seminal papers by Pollak and Wales (1979) and Fisher (1987; 1990) who find that welfare comparisons across households require consideration of both observed and unobserved differences. Although it is difficult to estimate models with unobserved heterogeneity when using utility functions, it is relatively straightforward to estimate statistical models of ‘well-being’ that incorporate observed and unobserved characteristics. A key barrier to using the latter models to explore this important issue further, however, is the inability to attach a welfare interpretation to the large body of results based on statistical well-being models. This paper overcomes this barrier by establishing a formal relationship between household utility and measures of well-being. The paper also derives measures of the welfare distortion attributable to unobserved household-level characteristics that can be estimated using models of well-being.

It is shown that there is a direct association between the equivalent income derived from well-being measures and the equivalent income derived from utility functions that define a household’s utility by reference to its specific subsistence threshold and its marginal utility of consumption. The conditions under which the two approaches generate isomorphic equivalent incomes are provided, including the relationship between the subsistence levels and marginal utilities of consumption defined by the household’s formal utility function and the household spending requirements and well-being (or welfare) sensitivities that can be deduced from models of well-being. The results provide a basis for determining welfare distortions using parameters obtained from typical estimable models of household well-being (eg. probit, logit and generalised extreme value models).

In order to investigate the welfare implications of unobserved heterogeneity, the properties of utility functions belonging to the Identical-Shape Hyperbolic-Absolute-Risk-Aversion (ISHARA) class are utilised in a novel manner such that heterogeneity is allowed in both subsistence levels and marginal utilities of consumption whilst continuing to accommodate exact linear aggregation across heterogeneous households. This is particularly useful since, as early as Gorman(1953), it has been known that the indifference

curves of (utilitarian) social welfare functions do not intersect if, and only if, there is exact linear aggregation. Given the sensitivity of welfare responses to general-equilibrium price changes caused by alternative policies, this implies that finding welfare-maximising policies is possible if and only if Gorman-type ISHARA utility functions are adopted (see, further, Gorman, 1961).

An Australian dataset on household-specific financial stress is used as a case study in order to empirically evaluate the welfare distortion measures derived in this paper. In evaluating these distortions, the proposed methodology considers the minimum level of household income required to compensate a household in a manner that accounts for its unique consumption requirements, and the shocks that the household has experienced (such as the illness or death of a family member).

In this respect, although the sources of household consumption heterogeneity have been explored (Clarida, 1991; Attanasio and Weber, 1995; Browning and Crossley, 2001; Gourinchas and Parker, 2002; Fernandez-Villaverde and Krueger, 2007), and the aggregate impact of such heterogeneity examined, significantly less is known about the distribution of unobserved subsistence levels and spending needs across households. Similarly, relatively little is known about the impact of such heterogeneity on the level of compensation households require in order to achieve a desired standard of living. The resulting level of compensation can be interpreted as a generalised equivalent income or a level of income that renders the household indifferent to its idiosyncratic habits, tastes and non-economic shocks. This is somewhat similar to the notion of an indifference scale (Chiappori, 2016).

The presence of unobserved heterogeneity implies a distribution of equivalent incomes for each household type. In this sense, the proposed approach complements the significant body of literature concerned with the estimation of models designed to capture unobserved heterogeneity (Briesch, Chintagunta and Matzkin, 2002 and 2009; Blundell, Kristensen and Matzkin, 2014; Lewbel and Pendakur, 2017). Under certain conditions, the moments associated with these distributions are shown to have direct welfare implications and are used to estimate, across a range of household types, the average welfare

distortion and the typical welfare gain or loss when there is a failure to appropriately account for heterogeneity.

The sources of these distortions are also examined, particularly in terms of whether the distortions are a function of heterogeneity in subsistence levels or heterogeneity in marginal utilities of consumption (van Praag, Goedhart and Kapteyn, 1980; Aiyagari, 1994; Calvet and Comon, 2003; Chetty and Szeidl, 2016). This distinction is important since distortions stemming from the former may be addressed through insurance markets or the elimination of frictions that produce adjustment costs (Blundell, Pistaferri and Preston, 2008; Lusardi, Schneider and Tufano, 2011). In contrast, the approach to correcting distortions associated with the latter is less clear.

The paper is organised as follows. Section 2 provides a formal relationship between measures of household utility and well-being, in addition to deriving measures of the distortion stemming from the assumption of homogeneity. Section 3 considers the estimation of the household-level expenditure requirements that are used to evaluate the measures proposed in the preceding section. The Australian dataset used in this paper is discussed in Section 4. Section 5 examines the welfare impact of group-wise heterogeneity, whereas Section 6 considers the welfare implications of deeper household-level heterogeneity. Concluding remarks are in Section 7.

2 A theorem for relating household well-being and utility functions

Consider an economy with N households. Each household has an ISHARA instantaneous utility function $u_{it}(c)$. The households can differ in their subsistence level such that household i 's utility function at time t is

$$u_{it}(c) = \frac{[ac + \beta_{it}]^{1-1/a} - 1}{a(1 - 1/a)}, \quad a > 0, a \neq 1, \beta_{it} < 0 \quad (1)$$

where $-\beta_{it}$ is household's i 's subsistence level of consumption at time t and a is a curvature parameter

The sign restrictions on a and β_{it} are motivated by cross-sectional evidence on savings rates (e.g., Dynan et al., 2004) that support the existence of subsistence levels of consumption and ISHARA utility functions with a positive common curvature parameter a . Pursuant to these restrictions, households can differ in their subsistence $-\beta_{it}$ such that any two households with the same level of consumption c and curvature parameter a will nevertheless attain a different utility level if their subsistence levels differ.

It can be shown that, if utility functions do not change over time, the household's utility function can be written as

$$u_i(c) = \exp(\zeta_i^*) \frac{[ac + \beta_i]^{1-1/a}}{a(1-1/a)} \quad (2)$$

where ζ_i^* is household i 's (relative) marginal utility of consumption (Mazzocco, 2007; Koulovatianos, Schroder and Schmidt, 2015). Consequently, household utility levels at a given consumption c may differ by reference to their subsistence level or their marginal utility of consumption.

Given (2), household j 's equivalent income takes on the general form

$$\bar{y}_j = \gamma_{ij} + \delta_{ij}\bar{y}_i \quad (3)$$

whereby \bar{y}_j is a linear function of household i 's income. This functional form is consistent with the well-known concept of Generalized Equivalence Scale Exactness (GESE) (Blackorby and Donaldson 1994; Donaldson and Pendakur, 2003; 2006; Cherchye et al., 2015).

To see that (3) holds, note that equivalising the utility of households i and j subject to (2) yields the equivalent income

$$\bar{y}_j = -a^{-1}(\beta_j - c_{ij}\beta_i) + c_{ij}\bar{y}_i \quad (4)$$

where $c_{ij} = \exp\left(-\frac{a}{a-1}(\zeta_j^* - \zeta_i^*)\right)$. Household j 's equivalent income is therefore a function of the subsistence levels of households i and j and differences in their marginal utilities of consumption.

Now consider the existence of a probability measure V_i that measures the household's 'well-being' and depends, at least partially, on household income. Note that the well-being function V_i can also depend on other factors and is not limited to income. Well-being is therefore defined generally and can be based on, for example, indicators of financial stress, credit default, financial satisfaction and even 'happiness'. The well-being function for household i takes on the form

$$V_i = V(h_i^y, x_i) = 1 - F\left(\frac{k_i(x_i) - y_i}{\eta_i}\right) \quad (5)$$

where $F(\cdot)$ is a cumulative distribution function (CDF), $k_i(x_i)$ is a household-specific income requirement or threshold that can depend on the set of characteristics x_i (for convenience $k_i(x_i)$ will be denoted k_i in the discussion that follows), y_i is household i 's income and η_i is a scale parameter. Since the change in $F(\cdot)$ depends on the size of the scale parameter η_i , we follow van Praag (1968; 1971) in labelling η_i as the household's 'welfare' (or well-being) sensitivity parameter. Technically, however, η_i reflects the household's welfare (in)sensitivity with $1/\eta_i$ representing well-being or welfare sensitivity.

The function $F(\cdot)$ is not chosen arbitrarily but on the basis that it encompasses the implicit well-being function adopted in much of the literature that investigates household-level well-being. In particular, the household's well-being is typically determined by reference to a linear probability model in one or more indicator variables that are explained by a set of characteristics x_i , which include household or individual income. An example is a model with a probit response variable that depends on household-level income and a set of (observed and unobserved) characteristics for each household.

It can be shown that the well-being or value function (5) has the following properties:

Definition 1. *Household i 's well-being or value function is defined and bounded between 0 and 1 if η_i is a strictly positive, bounded real number and \widehat{k}_i is a bounded, real number.¹*

¹Since household income is always a bounded real number.

Definition 2. Household i is balanced if its requirements are equal to its income, $k_i = y_i$, and its well-being function, if defined, is equal to 0.5.

Corollary 1. If the household's value function is defined it will approach zero monotonically in the threshold k_i and approach unity monotonically in household income y_i .

Corollary 2. As η_i approaches zero, household i 's value function becomes discrete and degenerates to the value 0.5 and the two polar extremes 0,1. The household's value function belongs to the set $\{0, 0.5, 1\}$ and any deviation from the balanced outcome renders the household completely 'satisfied' or completely 'unsatisfied'. The household's well-being is therefore a discrete, ordinal process that defines the household's preference ordering.

Corollary 3. As η_i approaches some arbitrarily large number, the value function is always 0.5.

Corollary 4. If only the outcome $V_i=0.5$ and household income y_{it} are observed, it is not possible to distinguish between a balanced household and a household with an arbitrarily large well-being (in)sensitivity term η_i .

Given (5), the solution to the well-being equivalence problem (whereby household j 's income y_j is adjusted so that its well-being function V_j is equal to household i 's well-being V_i)

$$\min_{y_j} \|V(\bar{y}_i, x_i) - V(y_j, x_j)\| \quad (6)$$

yields the following 'equivalent income' \bar{y}_j

$$\bar{y}_j = k_j + \frac{\eta_j}{\eta_i} (\bar{y}_i - k_i). \quad (7)$$

It should be noted that \bar{y}_j is the income level that equivalizes the well-being functions V_i and V_j rather than the utility functions u_i and u_j . The equivalent income \bar{y}_j for household j is, therefore, a linear function of the reference household's 'excess' income $y_i - k_i$, with an intercept given by its threshold k_j and a slope parameter that depends on the ratio of scale terms $\frac{\eta_j}{\eta_i}$.

Setting $\gamma_{ij} = k_j - \frac{\eta_j}{\eta_i} k_i$, $\delta_{ij} = \frac{\eta_j}{\eta_i}$, where (k_i, η_i) and (k_j, η_j) are the income requirement and scale parameters stemming from the well-being functions for households i and j respectively, we obtain

$$\bar{y}_j = \left(k_j - \frac{\eta_j}{\eta_i} k_i \right) + \frac{\eta_j}{\eta_i} \bar{y}_i \quad (8)$$

where it is clear that (8) accords with (4) and (3).

Theorem 1. *The equivalent incomes generated by the utility function $u_{it}(c)$ and the value function V_{it} are isomorphic and imply that household j 's equivalent income is a linear function of household i 's income. A one-to-one mapping between the parameters of the two functions can be established.*

The relationship between the two equivalent incomes is given by

$$\frac{\eta_j}{\eta_i} = c_{ij} = \exp \left(-\frac{a}{a-1} (\zeta_j^* - \zeta_i^*) \right) \quad (9a)$$

$$k_j = -a^{-1} \beta_j \quad (9b)$$

$$k_i = -a^{-1} \beta_i. \quad (9c)$$

Equation (9a) identifies the link between household i 's welfare sensitivity η_i and the marginal utility of consumption ζ_i^* . It is clear that in the case $\zeta_i^* = \zeta_j^*$, such that marginal utilities of consumption are the same for households i and j , we also have $c_{ij} = \frac{\eta_j}{\eta_i} = 1$. In this case, the same law of motion $\partial \bar{y}_j = \partial \bar{y}_i$ holds for both a household that maximizes V_i or its utility (2).

The remaining two equations (9b) and (9c) identify the link between the household's subsistence $-\beta_i$ in (2) and the household's income requirement k_i . The relationships imply that the difference in income requirements $k_j - k_i$ is proportional to $\beta_j - \beta_i$, with the proportionality depending on the curvature parameter a . A common proportionality term $(-a^{-1})$ implies a representative household (in other words, linear aggregation), although the more general equivalent income (8) is not bound by this restriction.

The results establish the following two corollaries which associate the parameters from CDF-based models of well-being to those stemming from a formal utility function.

Corollary 5. *If $a > 1$, the log of the ratio of welfare sensitivities for households i and j , $\ln \frac{\eta_j}{\eta_i}$, is a negative affine function of the difference in marginal utilities of consumption $\zeta_j^* - \zeta_i^*$.*

In order to better understand this, it is useful to explore the situation where $\ln \eta_j = \zeta_j^* = 0$ and $a > 1$. In this scenario, we have $\ln \eta_i = -\frac{a}{a-1}\zeta_i^*$ such that the log scale or welfare (in)sensitivity η_i is negatively related to the household's marginal utility of consumption.

Corollary 6. *Household i 's income requirement k_i is a positive affine function of its subsistence level $-\beta_i$.*

2.1 Determining welfare distortions stemming from the assumption of homogeneity

To examine the impact of homogeneity constraints assume that k_j is a linear function of a fixed component γ_{0j} and observed household characteristics x_j such that

$$\widehat{k}_j = \gamma_{0j} + \widetilde{c}_j = \gamma_{0j} + x_j' \gamma. \quad (10)$$

where $\widetilde{c}_j = x_j' \gamma$ for notational convenience, γ_{0j} is unknown, and \widehat{k}_j is the expected value of k_j . Note that \widehat{k}_j is inferred from the well-being function V_j .

For a given level of household income, households are therefore heterogeneous in terms of their unique spending requirement γ_{0j} , the scale parameter η_j and household characteristics \widetilde{c}_j of which the first two are entirely unobserved. From (9a)-(9c), this implies heterogeneity in the household's subsistence level β_j and marginal utility of consumption ζ_j^* .

Relative to a benchmark household characterised by the triple $(y^*, \widehat{k}^*, \eta^*)$, household j 's welfare-equivalising income and associated equivalence scale are

$$\bar{y}_j = \widehat{k}_j + \frac{\eta_j}{\eta^*} (y^* - \widehat{k}^*) = \gamma_{0j} + \widetilde{c}_j + \frac{\eta_j}{\eta^*} (y^* - \gamma_0^* - \widetilde{c}^*) \quad (11)$$

$$\frac{\bar{y}_j}{y^*} = \frac{\eta_j}{\eta^*} + \frac{\widehat{k}_j - \frac{\eta_j}{\eta^*} \widehat{k}^*}{y^*}. \quad (12)$$

To determine the impact of imposing homogeneity on household consumption preferences, consider the equivalence scale stemming from the assumption that $\eta_j = \eta^*$ and/or $\gamma_{0j} = \gamma_0^*$. These restrictions yield a homogeneous form of equivalent income, and the difference between (11) and the homogeneous measure is the additional income required by household j to achieve the same level of well-being as the benchmark household after taking into account its specific subsistence requirement and/or marginal utility of consumption. We consider three cases: (1) homogeneity in the marginal utility of consumption, $\eta_j = \eta^*$; (2) homogeneity in the fixed part of the subsistence level, $\gamma_{0j} = \gamma_0^*$; and (3) both forms of homogeneity.

Case 1: Homogeneity in the marginal utilities of consumption

If it is first assumed that marginal utilities of consumption are homogeneous such that η_j is equal to η^* then equivalent income, denoted as $\bar{y}_j | (\eta_j = \eta^*)$, is given by household income for the benchmark household y^* adjusted by the difference in the household-specific income needs of the two households

$$\bar{y}_j | (\eta_j = \eta^*) = y^* + (\widehat{k}_j - \widehat{k}^*) \quad (13)$$

such that the equivalence scale (12) simplifies to

$$\frac{\bar{y}_j}{y^*} | (\eta_j = \eta^*) = 1 + \frac{\widehat{k}_j - \widehat{k}^*}{y^*}. \quad (14)$$

It can be shown that the distortion $D_{\eta,j}$ stemming from common marginal utilities of consumption is non-zero unless the true η_j is equal to η^*

$$D_{\eta,j} = \bar{y}_j - \bar{y}_j | (\eta_j = \eta^*) = \left(\frac{\eta_j}{\eta^*} - 1 \right) (y^* - \widehat{k}^*) \quad (15)$$

where $D_{\eta,j}$ is defined as the difference between heterogeneous equivalent income \bar{y}_j and homogeneous equivalent income $\bar{y}_j | (\eta_j = \eta^*)$.

In turn, the bias in the equivalence scale stemming from the assumption that $\eta_j = \eta^*$ is a linear function of \widehat{k}^* with intercept $\left(\frac{\eta_j}{\eta^*} - 1\right)$ and slope parameter $\left(\frac{1-\eta_j/\eta^*}{y^*}\right)$

$$D_{\eta,j}^* = \frac{\bar{y}_j}{y^*} - \frac{\bar{y}_j}{y^*} \Big|_{(\eta_j = \eta^*)} = \left(\frac{\eta_j}{\eta^*} - 1\right) + \left(\frac{1 - \eta_j/\eta^*}{y^*}\right) \widehat{k}^*. \quad (16)$$

Since η_j is non-negative and unbounded on the right, the bias stemming from the adoption of $\eta_j = \eta^*$ lies in the interval $(\widehat{k}^* - y^*, \infty)$. Assuming $y^* > \widehat{k}^*$, this implies that $\frac{\bar{y}_j}{y^*} \Big|_{(\eta_j = \eta^*)}$ may constitute either an over- or under- estimate of the household income required to equate household j 's welfare with that of the benchmark household. As the interval is unbounded on the right, however, there is a greater risk that $\frac{\bar{y}_j}{y^*} \Big|_{(\eta_j = \eta^*)}$ understates the additional compensation required to equalise welfare for household j .

Figure 1 shows the bias in the equivalence scale when the true ratio $\frac{\eta_j}{\eta^*}$ is equal to 2 for a benchmark household with $y^* = \$75,000$. At $\widehat{k}^* = \$50,000$, for example, the assumption $\eta_j = \eta^*$ understates the equivalence scale by 1/3, with household j requiring an additional $1/3 \times y^* = \$25,000$ to equalise welfare in the presence of differences in the marginal utility of consumption.

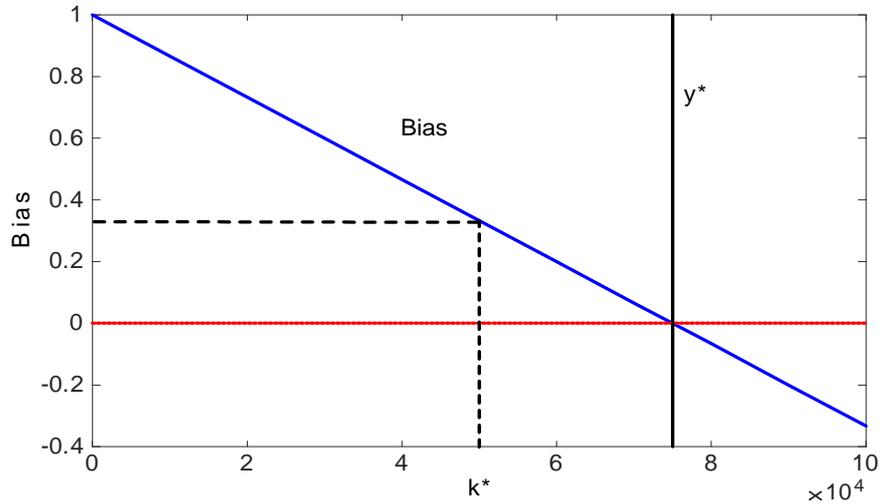


Figure 1: Bias in the equivalence scale given the assumption $\eta_j = \eta^*$ in the case where the true η_j/η^* is equal to 2. h^* is the benchmark household's income, with k^* representing the benchmark household's income requirement.

Case 2: Homogeneity in subsistence levels

Pursuant to (10), the household's subsistence level is decomposed into the part explained by its observed characteristics \tilde{c}_j and the unobserved part that is unique to the household γ_{0j} . If it is assumed that there is no household-specific component, such that $\gamma_{0j} = \gamma_0^*$, then households will differ in their subsistence levels only by reference to their known characteristics (for example, due to age differences in children). In this case, we obtain

$$\bar{y}_j | (\gamma_{0j} = \gamma_0^*) = \hat{k}_j + (\gamma_0^* - \gamma_{0j}) + \frac{\eta_j}{\eta^*} (y^* - \hat{k}^*) \quad (17)$$

$$\frac{\bar{y}_j}{y^*} | (\gamma_{0j} = \gamma_0^*) = \frac{\eta_j}{\eta^*} + \frac{1}{y^*} \left(\hat{k}_j - \frac{\eta_j}{\eta^*} \hat{k}^* \right) + \frac{1}{y^*} (\gamma_0^* - \gamma_{0j}). \quad (18)$$

Consequently, even if the assumption of common marginal utilities of consumption is relaxed, there may be distortion stemming from the assumption of a common subsistence level.

It can be shown that this distortion is an affine function of the simple difference $\gamma_{0j} - \gamma_0^*$ and will therefore only be zero in the case where the true γ_{0j} is equal to γ_0^*

$$D_{\gamma_{0j}} = \bar{y}_j - \bar{y}_j | (\gamma_{0j} = \gamma_0^*) = \gamma_{0j} - \gamma_0^* \quad (19)$$

$$D_{\gamma_{0j}}^* = \frac{\bar{y}_j}{y^*} - \frac{\bar{y}_j}{y^*} | (\gamma_{0j} = \gamma_0^*) = \frac{1}{y^*} (\gamma_{0j} - \gamma_0^*). \quad (20)$$

where $D_{\gamma_{0j}}$ is the additional income that would be required to equalize household j 's well-being with that of the benchmark household (in the case where $\gamma_{0j} = \gamma_0^*$ were incorrectly assumed). $D_{\gamma_{0j}}^*$ is the distortion in the equivalence scale.

Case 3: Homogeneity in both marginal utilities of consumption and subsistence levels

Finally, it can be shown that the distortions when both $\eta_j = \eta^*$ and $\gamma_{0j} = \gamma_0^*$ are imposed are

$$\begin{aligned} D_j &= \bar{y}_j - \bar{y}_j | (\eta_j = \eta^*, \gamma_{0j} = \gamma_0^*) \\ &= \left(\gamma_{0j} - \frac{\eta_j}{\eta^*} \gamma_0^* \right) + \left(1 - \frac{\eta_j}{\eta^*} \right) \tilde{c}^* + \left(\frac{\eta_j}{\eta^*} - 1 \right) y^* \end{aligned} \quad (21)$$

$$\begin{aligned}
D_j^* &= \frac{\bar{y}_j}{y^*} - \frac{\bar{y}_j}{y^*} | (\eta_j = \eta^*, \gamma_{0j} = \gamma_0^*) \\
&= \left(\frac{\eta_j}{\eta^*} - 1 \right) + \frac{1}{y^*} \left[\gamma_{0j} - \frac{\eta_j}{\eta^*} \gamma_0^* + \left(1 - \frac{\eta_j}{\eta^*} \right) \tilde{c}^* \right].
\end{aligned} \tag{22}$$

Equations (21) and (22) imply that the assumption of homogeneity in the η_j and γ_{0j} parameters may over- or under- state equivalent income. Ceteris paribus, however, if the household has a greater subsistence requirement, such that $\gamma_{0j} > \gamma_0^*$, then the assumption of homogeneity will *under*-state the household's equivalent income resulting in a reduced level of well-being. Conversely, $\gamma_{0j} < \gamma_0^*$ will *over*-state equivalent income resulting in a greater level of well-being.

Consider also the case where $\eta_j < \eta^*$, such that household j exhibits a greater level of welfare sensitivity to a change in income than the benchmark household. In this case, the assumption of homogeneity will *over*-state the household's equivalent income (with the converse also holding true). The reason for this result is that η_j implies a smaller income shift to equalise $V(y_j, x_j)$ with $V(y^*, x^*)$ than does η^* . In particular, when $\eta_j < \eta^*$, the slope of household j 's marginal utility of consumption is flatter if η^* is imposed instead of η_j . As such, achieving the absolute welfare change $|\Delta V(y_j, x_j)|$ requires a greater change in household income y_j than if the true η_j were adopted. The converse holds if $\eta_j > \eta^*$.

3 Modelling household income requirements

It is clear from the preceding discussion that in the absence of exogenous values for β_i and ζ_i^* , the mapping of parameters from V_i to u_i is not unique and will depend on the policy-maker's choice of V_i . Given the plausible relationship between the household's capacity to consume up to its subsistence level and its experience of financial stress, V_i is chosen such that it reflects estimates of expenditure requirements that are based on the avoidance of financial stress. The econometrician is, however, free to use other models in order to underpin V_i (consider, for example, the models in Senik (2004), Zaidi and Burchardt (2005), Morciano, Hancock and Pudney (2015) or Decancq, Fleurbaey and Schokkaert (2017)).

Pursuant to the chosen well-being function, households are averse to a range of indicators of financial stress (for example, needing to borrow money from friends or family or being unable to pay utility bills) and prefer a lower probability of financial stress. Under (5) or (2), households seek to maximize the distance between their income and some minimum commitment level.

The model requires estimates of k_{it} across households $i = 1, 2, \dots, N$. This raises the issue of the suitable identification of minimum expenditure or commitment levels. To achieve identification, binary indicators of household-specific financial stress are adopted. Consequently, the household is simply required to acknowledge its engagement in some activity that reflects financial stress and need not explicitly consider its minimum expenditure level, which is prone to bias (de Ree, Alessie and Pradhan, 2013).

To formalise the financial stress condition, first define the indicator variable m_{it} as equal to 1 when household i is financially stressed in period t and 0 otherwise. The indicator is constructed using

$$m_{it} = I(k_{it} > r_{it}) \quad (23)$$

where k_{it} is the household's unobserved income requirement, r_{it} is the household's actual income net of its actual accommodation expenditure (hereafter called the household's residual income) and $I(\cdot)$ is a binary indicator taking on the value unity if $k_{it} > r_{it}$.

Without loss of generality, r_{it} is adopted instead of y_{it} . For example, if the household's accommodation expenditure is unobserved the decision rule becomes $I(k_{it}^+ > y_{it})$, with k_{it}^+ now inclusive of housing expenditure. If housing costs are observed, however, it is preferable that they be concentrated out of the household's decision rule since housing accommodation is almost certainly part of the household's required spending. This also improves the capacity to isolate the habits and tastes underpinning the household's non-housing consumption needs.

The existence of financial stress depends on (but is not limited to) hierarchical criteria such as whether the household has pawned or sold something to make ends meet, required financial assistance from family or friends, could not pay utility bills on time or could not afford to heat its home. For the case study adopted in this paper, the financial stress

indicators in the Household Income and Labour Dynamics in Australia (HILDA) survey are adopted (see, further, Section 4.1).

Given the hierarchical nature of the identification regime, it is instructive to consider the interpretation of the group-wise minimum expenditure difference $E [k_t^{l^*} - k_t^l]$. This value represents the difference between the (non-housing) spending requirements of the households belonging to groups l^* and l in period t (to avoid confusion, the use of the super- or sub- script l hereafter refers to a group of households whereas i, j refer to individual households). Assume that households A and B are in a similar financial position, belong to groups l^* and l respectively, and that only household A has access to friends or family that it seeks financial assistance from. Household A obtains financial assistance and is therefore deemed to be financially stressed. Since household B is unable to access assistance from friends or family, its financial stress must be captured by other indicators. For example, assume that in the absence of financial assistance household B is placed in a position where it is unable to pay its utility bills on time; in this case its financial stress is captured pursuant to the inherent hierarchy in the stress criteria.

Assume, however, that household B 's stress is not identified by any of the stress criteria. Household B is therefore deemed not to be in a position of financial stress solely due to its inability to find family or friends willing to lend to it. On the basis that constraints on borrowing from family or friends across groups l^* and l are randomly selected from a common distribution (or from two distributions with a common location), the estimate $E [k_t^{l^*} - k_t^l]$ will provide the difference between the expected minimum expenditure requirements of the two groups notwithstanding the presence of borderline circumstances such as the one aforementioned.

Parameterizing the household's income requirement k_{it}

Assume that the household forms an expectation of its spending commitments and therefore knows its expected income requirement \widehat{k}_{it} . Since the household's spending needs are a function of the unknown (to the econometrician) consumption commitments and adjustment costs of its individual members, each household has its own household specific spending requirement γ_{0i} that jointly characterises the needs of its individual

members (the importance of adjustment costs is discussed in Chetty and Szeidl, 2007; 2016). When the household’s composition exhibits a substantive change, the household is allocated a new γ_{0i} in order to reflect its ‘updated’ composition (see, further, Section 4.2).

Household i ’s estimated income requirement is given by

$$\widehat{k}_{it} = \gamma_{0i} + x'_{it}\gamma \quad (24)$$

where x_{it} is a set of covariates or instruments used to estimate time-variation in household i ’s spending needs (refer to the Appendix). These include, for example, household specific events such as changing residency, major events such as serious illness, and year-specific macro effects that are common to all households.

The solution \bar{y}_i to the problem (6) yields the income required by household i to achieve the benchmark welfare level $\bar{V} = V(y^*, x^*)$. Given the specification for \widehat{k}_{it} , this income level differs from the traditional equivalent income measure as \bar{y}_i is effectively the level of income that renders a household *indifferent* (in terms of its financial stress) to its idiosyncratic consumption and welfare sensitivity differences. The resulting \bar{y}_i can be termed either an equivalent income or an indifference income although the former term is adopted here (see, also, Browning, Chiappori and Lewbel, 2013). The associated ‘equivalence’ scale is expressed as the ratio \bar{y}_i/y^* .

Since the econometrician is unable to observe k_{it} , an error term u_{it} is introduced such that

$$k_{it} = \widehat{k}_{it} + u_{it} \quad (25)$$

$$= \gamma_{0i} + x'_{it}\gamma + u_{it} \quad (26)$$

$$u_{it} \sim N(0, \eta_i^2). \quad (27)$$

The interpretation of η_i is somewhat convoluted as it encompasses the econometrician’s uncertainty about k_{it} , measurement error and the weighted average of the welfare sensitivities of the individuals in the household (see, further van Praag (1968; 1971) and

Goedhart et al. (1977)). Heterogeneity in welfare sensitivity reflects a range of factors such as the household's tastes and habits, and its capacity to adjust its spending in response to an income shock. As with γ_{0i} , the value of η_i is allowed to change following a substantive change in the household's individual composition.

Determining the particular proportion of η_i that is attributable to the econometrician's uncertainty rather than welfare sensitivity is outside the scope of this paper. However, if one is prepared to assume that the econometrician's uncertainty regarding k_{it} is similar across households (viz. the econometrician does not have greater knowledge of household i 's income requirement k_{it} than household j 's requirement k_{jt}), then heterogeneity in η_i is interpretable in terms of differences in welfare sensitivity. Indeed, Section 4.3 presents evidence that strongly supports this interpretation.

Given (26) and (27), household i 's probability of financial stress is a function of its residual income r_{it} , the household-specific component of its subsistence level γ_{0i} , its welfare sensitivity (which reflects its marginal utility of consumption) η_i , and the information incorporated in the observed regressors x_{it} . It can be shown that this probability is given by

$$\Phi\left(\frac{\gamma_{0i} + x'_{it}\gamma - r_{it}}{\eta_i}\right) = \Phi\left(\frac{\gamma_{0i} + x'_{it}\gamma + \bar{h}_{it} - y_{it}}{\eta_i}\right) \quad (28)$$

where $\Phi(\cdot)$ is the standard normal distribution function and \bar{h}_{it} is the household's housing accommodation expenditure. It is clear from (28) that a greater η_i renders the household less sensitive to a change in income y_{it} , whereas low values of η_i will produce sharp changes in the household's well-being function following a change in income.

The resulting well-being function depends on the extent to which household i 's income is sufficient to meet its expenditure requirements (which, in turn, reflect the household's subsistence level)

$$V_{it} = 1 - \Phi\left(\frac{\widehat{k}_{it} - r_{it}}{\eta_i}\right) = 1 - \Phi\left(\frac{\gamma_{0i} + x'_{it}\gamma + \bar{h}_{it} - y_{it}}{\eta_i}\right) \quad (29)$$

with V_{it} clearly according with (5) such that the welfare interpretations for k_{it} and η_i stemming from Theorem 1 hold.

4 Data and estimation

4.1 Data

The HILDA dataset is used as a case study for examining the welfare distortions stemming from unobserved heterogeneity. The dataset is based on a survey which is undertaken annually and collects information about economic and subjective well-being, labour market dynamics and family dynamics in Australian households. The first wave (in 2001) of the survey consisted of 7,682 households, with an additional 2,153 households added in 2011 (viz. wave 11). For the purposes of this application, the dataset's main limitation is that the number of observations per household is limited which affects the precision of the parameter estimates for $\hat{\gamma}_{0i}$ and $\hat{\eta}_i$. This is a problem that is generally observed in large household level datasets. This limitation, however, does not appear to affect the general conclusions drawn from the case study and this is discussed further in Sections 4.3 and 5.

Given the model's dependence on household income levels, households are identified by reference to the main income earner. If individual i is identified as the main income earner in waves t , $t + 1$ and $t + 2$ then the set of information (personal and household) associated with individual i over that period, $\{r_{it}, r_{it+1}, r_{it+2}, x_{it}, x_{it+1}, x_{it+2}\}$ is identified as belonging to a single household.

Residual income r_{it} is constructed as the difference between household income and housing accommodation costs which are based on the household's annual mortgage repayments or housing rent. A very small number of households made both rental and mortgage repayments. In this case, the sum of both values is treated as their housing-related expenditure. Without loss of generality, the model is estimated after dividing r_{it} by $fs_{it} \times 5000$, where fs_{it} is household size in period t , such that the per-person threshold $k_{it}/(fs_{it} \times 5000)$ is estimated in bundles of \$5000. Residual income r_{it} is also converted to real terms such that k_{it} represents real spending needs and any time effects in k_{it} are not the result of inflation. To avoid excessive notation, in the remaining sections of the paper any reference to k_{it} or r_{it} is on a real per-person basis except where stated otherwise.

The indicator m_{it} is constructed based on household responses to a set of financial stress variables in HILDA. In particular, it is assumed that a household is financially stressed (viz. $m_{it} = 1$) if it cannot pay its utility bills, mortgage or rent on time, requested financial help from friends or family, pawned or sold something to make ends meet, is unable to heat its home, goes without meals or requests help from a charity or similar organisation.

4.2 Household formation

Households are dynamic and the dataset is not limited to households with constant characteristics. This is critical since changes to demographic characteristics are needed to identify household preferences regarding children, fertility and demographic composition. Households can also change location, or may be renters in period t and home owners in period $t + 1$. However, if the primary wage earner changes in, for example, period $t + 3$ then this is treated as a ‘new’ household which is allowed to have its own γ_{0i} and η_i parameters. As such, the estimated values of γ_{0i}, η_i are sensitive to significant changes in the household’s income structure. This process is iterated commencing at wave 2 and ending at wave 12 of the survey, thereby identifying households h_i and their corresponding dataset $\{m_{it}, m_{it+1}, \dots, r_{it}, r_{it+1}, \dots, x_{it}, x_{it+1}, \dots\}$.

To enable estimation, the sample is restricted to households for which at least one switch in the indicator variable takes place. Observations are also removed when it is not possible to deduce household income, financial stress or when x_{it} is not observed. Following Krueger and Perri (2006), only households that have been interviewed at least 6 times are considered (however, changing this value to 8 or 10 produces similar results).

In order to assess the robustness of the formations process, parameter estimates are also obtained subject to the identification of households using the first survey respondent rather than the main income earner and by limiting the dataset to households that do not change. The alternative identification methods, however, do not produce large changes in the parameters or the estimated distortions.

4.3 Estimation and model diagnostics

Each household is characterised by the set of observables $\{m_i, X_i, r_i\}$ for $i = 1, 2, \dots, N$. X_i is comprised of x_{it} for $t = t_{0i}, t_{0i} + 1, \dots, T_i$ (where t_{0i}, T_i index the start and end periods of the i th household). The i th household's contribution to the likelihood function L is

$$L_i(m_i|\gamma, \gamma_{0i}, \eta_i, X_i) = \prod_{t=t_{0i}}^{T_i} \Phi\left(\frac{\gamma_{0i} + x'_{it}\gamma - r_{it}}{\eta_i}\right)^{m_{it}} \left(1 - \Phi\left(\frac{\gamma_{0i} + x'_{it}\gamma - r_{it}}{\eta_i}\right)\right)^{1-m_{it}} \quad (30)$$

with the model's overall likelihood function given by

$$L(m|\gamma, \gamma_0, \eta, X) = \prod_{i=1}^N L_i(m_i|\gamma, \gamma_{0i}, \eta_i, X_i). \quad (31)$$

In line with much of the literature in this space, the model contains a large number of parameters that are difficult to estimate. To alleviate this problem, a Bayesian approach to learning about the model parameters is adopted along the lines of Buera, Monge-Naranjo and Primiceri (2011) who encounter similar considerations using a model with a likelihood function analogous to that adopted in this paper. In particular, prior distributions for \widehat{k}_{it} and η_i^2 are specified and the model parameters are chosen by maximising the resulting posterior density formed by the product of the likelihood function and the prior densities. Note that a prior distribution is placed on \widehat{k}_{it} rather than its constituent parameters (γ_{0i}, γ) . The reason for this is two-fold. First, there is a clearer prior on the value of \widehat{k}_{it} than on the values of its constituent parameters. Second, parameter estimates of the individual parameters (γ_{0i}, γ) are obtained that are not directly affected by the choice of prior. The chosen priors are

$$\widehat{k}_{it} \sim N(\mu_0, \sigma_0^2) \quad (32)$$

$$\eta_i^2|\widehat{k}_{it} \sim IG(n_0, s_0) \quad (33)$$

where $N(\cdot)$ is the normal density with location and scale parameters μ_0, σ_0 , and IG is

the Inverse-Gamma density with shape n_0 and scale s_0 .

Fortunately, the priors for \widehat{k}_{it} and η_i^2 have clear interpretations with μ_0 representing prior beliefs regarding the typical household's spending requirements and σ_0 reflecting uncertainty about the variation in this value. Similarly, the hyper-parameters (n_0, s_0) reflect (conditional on \widehat{k}_{it}) the household's probability of being financially stressed and the uncertainty regarding this probability.

Importantly, although the priors influence the unconditional values of \widehat{k}_{it} and η_i^2 they impose little restriction on the cross-household variation in these parameters which this paper is primarily interested in. This is readily evident in the paper's case study with the distribution of \widehat{k}_{it} across households clearly being non-normal.

The prior hyper-parameters k_{it} are set to $\mu_0 = 23,000/5000, \sigma_0 = 4$ such that the prior assumption is that each person requires expenditure of approximately \$23,000 per annum to avoid stress. This figure is based on the average residual income for households when $m_{it} = 1$. A large standard deviation is adopted for the prior to reflect the uncertainty regarding the expected value of k_{it} . As such, the prior chosen is not particularly informative.

The hyper-parameters (n_0, s_0) are set such that η_i has a prior mean of unity and a standard deviation of 0.5. These values imply that, given a typical level of residual income per person, the unconditional probability of a household being financially stressed is about 1 in 3 (with a 90 per cent chance of being between 15 and 41 per cent). The prior is, therefore, relatively uninformative and reflects the overall financial stress levels reported in Australian Bureau of Statistics' (ABS) data on household poverty (ABS, 2004; 2011). The prior is also consistent with Lusardi, Schneider and Tufano (2011) who find that around 25 per cent of US households are financially fragile.

The model is estimated on an unbalanced panel containing 26,397 observations of the indicator variable m_{it} (with 9,608 instances where $m_{it} = 1$) covering $N = 3,103$ households. The estimation process uses an analytical gradient and Hessian. The sparseness of the Hessian is exploited to obtain its inverse, and the posterior mode is retrieved using a Newton-Raphson algorithm. Given the analytical gradient and Hessian, convergence is

fairly stable given sensible starting conditions.

To examine the impact of household lifespan on the parameter estimates, the model is also estimated by restricting the lifespan to an initial 6 periods for all households and then successively re-estimating by adding an additional period's worth of data each time. This does not have a significant impact on the estimates of γ_{0i} or η_i and, most importantly, has little impact on the cross-sectional variation in these parameters that this paper focuses on.

The model appears to perform well in identifying stressed households, predicting the financial stress of the majority of households and nearly doubling the basic probit (or logit) model's capacity to predict financial stress. The results provide empirical support for the proposition that a household's propensity to experience financial stress is related to its capacity to meet its expenditure requirement \hat{k}_{it} .

Table 1 provides diagnostics for the estimated model which are based on the posterior mode of the model's parameters and contrasts the diagnostics with those from a standard probit model (a logit model was also estimated and performed similarly to the probit model).² The model is able to clearly distinguish between financially stressed and non-stressed households with the average predicted outcome during periods of stress, $E(\hat{m}_{it}|m_{it} = 1)$, being 58 per cent and the average predicted outcome during periods without stress, $E(\hat{m}_{it}|m_{it} = 0)$, being just under 28 per cent. As a point of comparison, the standard probit model has difficulty distinguishing between financially stressed and non-stressed households with the difference between the expected conditional probabilities $E(\hat{m}_{it}|m_{it} = 1) - E(\hat{m}_{it}|m_{it} = 0)$, being just under 13 per cent. In terms of prediction accuracy, the model is able to predict 62 per cent of financial stress episodes, whereas the probit model correctly predicts only 35 per cent of such episodes.

The proposed model also provides an accurate depiction of the dispersion of financial stress, with predicted levels of stress being close to actual levels. In particular, the model predicts financial stress 37.3 per cent of the time, which is close to the actual value of 36.4 per cent thereby resulting in a bias of less than 1 per cent. The probit model, on the other

²The individual parameter estimates for β, γ used to obtain k_{it} are typically statistically significant at the 0.05 level and are available on request.

hand, clearly under-predicts with only 20.1 per cent of observations being associated with financial stress thereby inducing a negative bias of 16.3 per cent.

Table 1: Model diagnostics

	Proposed model	Standard probit
<i>Obs</i>	26,397	26,397
$E(\hat{m}_{it} m_{it} = 0)$	0.277	0.323
$E(\hat{m}_{it} m_{it} = 1)$	0.582	0.453
$E(\hat{m}_{it} > 0.5 m_{it} = 1)$	0.616	0.354
$E(\hat{m}_{it} \leq 0.5 m_{it} = 0)$	0.764	0.886
<i>Bias</i> [#]	0.009	-0.163

Note: The bias is computed as the difference between the predicted and actual stress incidence, $E(\hat{m}_{it} > 0.5) - E(m_{it})$.

The estimated welfare sensitivities and the welfare sensitivity curve

Another issue that is of interest is the extent to which $\hat{\eta}_i$ incorporates information about household welfare sensitivities rather than the econometrician's uncertainty or measurement error. Theory suggests that welfare sensitivity should decline with income, particularly since higher income households are better able to smooth their consumption levels (Krueger and Perri, 2006; Blundell, Pistaferri and Preston, 2008). It follows, therefore, that $\hat{\eta}_i$ should increase as household income rises (thereby reducing sensitivity to a shift in income). Another way to understand this is based on Corollary 5, whereby $\hat{\eta}_i$ can be formally interpreted in terms of the reciprocal of the household's marginal utility of consumption. In this situation, theory suggests that $1/\hat{\eta}_i$ should decline as household income rises.

In the absence of any information about welfare sensitivity in the likelihood function, however, the adoption of an equal prior on η_i for all households would produce estimates $\hat{\eta}_i$ that are invariant to household income. Furthermore, if $\hat{\eta}_i$ only reflected the econometrician's uncertainty about the household's spending needs there would be no a priori reason to observe any systematic positive correlation between $\hat{\eta}_i$ and household income.

Figure 2 shows $\hat{\eta}$ across household income strata. The figure identifies a particularly strong positive relationship between $\hat{\eta}$ and household income (alternatively, a strong negative relationship between the household's marginal utility of consumption and its income) that is clearly consistent with economic theory. The figure also produces, for the first time (to the author's knowledge), an empirical estimate of the welfare sensitivity curve.

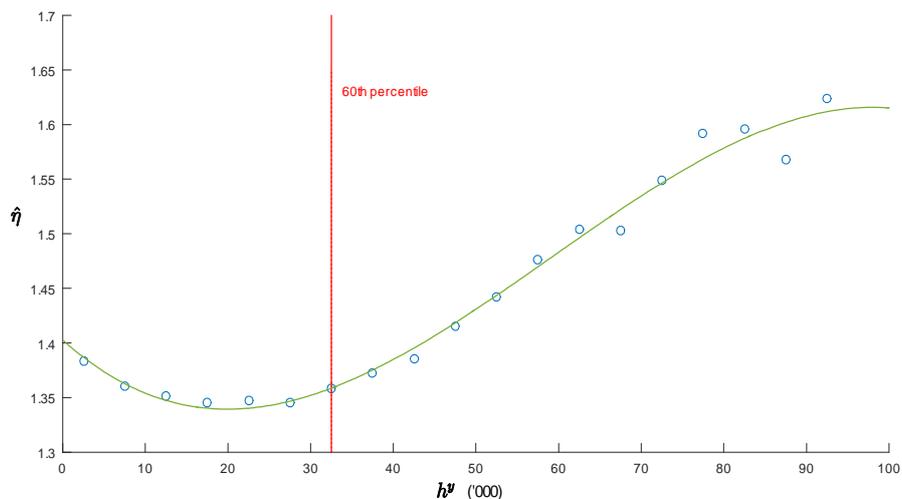


Figure 2. Welfare sensitivity curve. A decline in $\hat{\eta}$ implies greater welfare sensitivity. In accordance with Corollary 5, this can also be interpreted in terms of an increase in $1/\hat{\eta}$ implying a higher marginal utility of consumption. The curve is estimated as the sample mean of the welfare sensitivities of households sorted into per-capita household income (h^y) brackets of \$5000 (ending at \$95,000 - \$100,000). The vertical line indicates the 60th percentile of the distribution of per-capita household income. The fitted curve is a cubic polynomial.

This can be viewed as a generalised measure of the household's capacity to adjust its consumption following a change in its income. Although general evidence regarding welfare sensitivity is available, the empirical measurement of the full curve using unit-record data indicates clearly that the household's sensitivity to a change in income declines non-linearly with household income. It also indicates that the monotonicity in the relationship

between income and welfare sensitivity only seems to hold beyond a certain income level.

The results indicate fairly strongly that the estimates of $\widehat{\eta}_i$ are informative (about the marginal utility of consumption) in a manner that is consistent with economic theory and confirm the model's premise that households exhibit non-linear heterogeneity in both subsistence levels and in their marginal utility of consumption.

In particular, the shape of the curve suggests that welfare sensitivity actually rises at lower income levels before declining when income reaches a given threshold, and then plateauing for households in the right tail of the income distribution. This indicates that, although households are generally better able to adjust their consumption as their income rises, the capacity to adjust consumption initially declines at the lower income strata. Consequently, a rise in income for a *low income* household can result in a new consumption bundle that actually renders the household *more* sensitive to a change in income.

5 The impact of heterogeneity on group-wise welfare levels

The next two sections examine the distortions derived in Section 2. In this section, the paper evaluates whether there are meaningful *group-wise* differences in income requirements and welfare sensitivities, with a focus on demographic groups. Although this evaluation provides information about demographic-level differences, it does not consider the welfare impact of household level heterogeneity. Consequently, Section 6 examines the distribution of equivalence scales within each household type or group. Importantly, it examines the shape of the distribution of equivalence scales within each household type, and reports on the centrality, variability and skewness of the distributions.

Using various diagnostics, it is reasonably clear that the distribution of household specific income requirements and welfare sensitivities is not uniform across the various household types. The fixed component of income requirements $\widehat{\gamma}_{0l}$ is greatest for lone households, group households and households without children, and falls significantly in

the presence of children (Table 2). This is portrayed clearly when $\hat{\gamma}_{0l}$ is disaggregated by family size such that, with each additional family member, the fixed component of the household's per-capita income requirement becomes smaller. Welfare sensitivity also appears to differ by household type and size. In particular, $\hat{\eta}_l$ falls, hence welfare sensitivity increases, with larger family size. This implies that, although larger households have smaller per-capita fixed income requirements $\hat{\gamma}_{0l}$, they also have a smaller capacity to adjust consumption thereby resulting in greater sensitivity to income changes.

Table 2. Income requirement and welfare sensitivity parameters, $\hat{\gamma}_{0l}$ and $\hat{\eta}_l$, by household type and size

	Obs.	$\hat{\gamma}_{0l}$	$\hat{\eta}_l$
Household type			
Couple family w/out children	4237	3.717**	1.425**
Couple family with children < 15	4244	3.274	1.390
Couple family (dependent children)	513	3.449**	1.397
Couple family (non-dependent children)	641	3.716**	1.418**
Lone parent with children < 15	1432	3.090**	1.379**
Lone parent (dependent children)	284	3.539**	1.385
Lone parent (non-dependent children)	551	3.611**	1.408**
Lone person	5527	3.880**	1.428**
Group household	353	3.829**	1.422**
Multi family household	213	2.986**	1.376
Family size			
1	5527	3.979**	1.428**
2	5855	3.700**	1.419**
3	2904	3.546**	1.404**
4	2785	3.289	1.390
5	1127	2.868**	1.371**
6	370	2.808**	1.369**

Note: Values marked with ** indicate that the sample for the particular group is statistically different (at the .05 level) to the sample associated with the benchmark household (being a Couple family with children < 15 when grouping by household type or a family of size 4 when grouping by family size) using a Kolmogorov-Smirnov test.

Although the parameters in Table 2 are significantly different across groups in statistical terms, the welfare implications associated with these values are unclear. In this respect, consider a policy-maker who compensates household *groups* (as opposed to individual households which are examined in the section) with the objective of equalising group l 's welfare V_l with some benchmark \bar{V} . The policy-maker is required to make a choice regarding the extent to which group-specific requirements are accounted for in determining the level of any compensation. The welfare distortions associated with the assumptions depend on the biases (15) - (22). The biases will be zero if and only if the 'true' η_l and γ_{0l} parameters are equal to their benchmark equivalents η^* , γ_0^* . Otherwise, alternative assumptions regarding η_l and γ_{0l} may produce economically substantive differences.

The four alternative assumptions or scenarios considered in this analysis are:

(i) heterogeneity is allowed for the parameters reflecting both marginal utility of consumption η_l and the fixed part of subsistence γ_{0l} (whereby welfare differences are reflected in the equivalence scale \bar{h}_l^y/h_*^y);

(ii) heterogeneity is allowed only for γ_{0l} and group l 's marginal utility of consumption is forced to be equal to the benchmark η^* ($\bar{y}_l/y^*|\eta_l = \eta^*$);

(iii) heterogeneity is allowed only for η_l and the fixed part of group l 's subsistence is forced to be equal to the benchmark γ_0^* ($\bar{y}_l/y^*|\gamma_{0l} = \gamma_0^*$); and

(iv) homogeneity is forced for both the marginal utility of consumption and the fixed subsistence level ($\bar{y}_l/y^*|\eta_l = \eta^*, \gamma_{0l} = \gamma_0^*$).

To evaluate the welfare impact stemming from the failure to account for household-specific heterogeneity, the preceding four sets of equivalence scales are computed. The policy-maker determines welfare by choosing one of these four sets. The first set of equivalence scales is based on a policy maker that compensates for all group-level heterogeneity.

This policy maker chooses the equivalence scale given by the expected value of (12). The second and third sets of equivalence scales involve compensating only for differences in the marginal utility of consumption or the fixed part of the subsistence level and are based on the expected values of (14) and (18) respectively. The fourth and final set of equivalence scales is for a policy-maker who does not compensate for either γ_{0l} or η_l , but rather assumes that marginal utilities of consumption and fixed subsistence levels are the same for all groups.

The difference between the equivalence scales for scenarios (i) and (iv) therefore reflects the total distortion stemming from the failure to account for unobserved group-level heterogeneity. This distortion, and the equivalence scales for each of the four scenarios considered, are presented in Table 3. The benchmark household is assumed to be the average family of size 4 (in other words, the benchmark values $y^*, \eta^*, \gamma_0^*, \tilde{c}^*$ are sample averages across *all* households of size 4, including periods where these households were stressed) but choosing another type of reference household does not affect the general conclusion. A positive distortion indicates that the failure to account for differences in η and γ_0 results in *under*-compensation with the household requiring additional income to achieve the well-being of the benchmark household group (the converse holds for a negative distortion).

It is clear from Table 3 that the assumption of homogeneity will substantially over-compensate larger households, and under-compensate smaller (up to 3 person) households. In all cases, the distortions are statistically significant. A household with 5 members is typically over-compensated by about 12 percent if homogeneity is assumed, rising to nearly 17 per cent for a household with 6 members. In general, the relationship between distortion and family size implies that larger households typically require less to achieve the same welfare level suggesting that larger households, perhaps like older households (Aguiar and Hurst, 2005), engage in some form of substitution in order to increase their welfare.

The reason for this over-compensation is that, although larger households are less able to adjust their consumption bundles in response to an income change, they also

have smaller subsistence levels (on a per-person basis) and these *more* than offset their higher marginal utility of consumption.

Table 3: Equivalence scales under alternative assumptions regarding heterogeneity

	Heterog.	Homogeneous			Distortion
	(i)	(ii)	(iii)	(iv)	
		$\hat{\eta}_l = \eta^*$	$\hat{\gamma}_{0l} = \gamma_0^*$	both	
<i>Grouped by household type</i>					
Couple family w/out children	0.519	0.517	0.472	0.470	0.049**
Couple family with children < 15	1.000	1.000	1.000	1.000	N/A
Couple family (dependent children)	0.927	0.927	0.893	0.892	0.035**
Couple family (non-dep. children)	0.856	0.854	0.779	0.777	0.079**
Lone parent with children < 15	0.688	0.689	0.717	0.717	-0.029**
Lone parent (dependent children)	0.646	0.647	0.611	0.612	0.034**
Lone parent (non-dep. children)	0.549	0.548	0.511	0.510	0.039**
Lone person	0.270	0.269	0.238	0.237	0.033**
Group household	0.618	0.617	0.553	0.552	-0.066**
Multi family household	1.165	1.166	1.244	1.246	-0.081**
<i>Grouped by family size</i>					
1	0.285	0.285	0.248	0.247	0.038**
2	0.547	0.545	0.496	0.495	0.052**
3	0.806	0.806	0.759	0.758	0.048**
4	1.000	1.000	1.000	1.000	N/A
5	1.121	1.123	1.243	1.245	-0.124**
6	1.329	1.332	1.502	1.505	-0.176**

Note: Heterogeneous equivalence scales (column denoted *i*) use group *l*'s 'true' values $\hat{\eta}_l$ and $\hat{\gamma}_{0l}$ calculated as the sample averages of the parameters for households in group *l*. The homogeneous equivalence scales (columns denoted *ii,iii* and *iv*) assume that: welfare sensitivities are homogeneous ($\hat{\eta}_l = \eta^*$); the fixed part of the subsistence level is homogeneous ($\hat{\gamma}_{0l} = \gamma_0^*$); or both $\hat{\eta}_l$ and $\hat{\gamma}_{0l}$ are homogeneous.

The estimated distortion is the difference between the equivalence scales obtained when assuming full heterogeneity and full homogeneity (columns denoted i and iw respectively). ** indicates significance at the .01 level.

A comparison of the equivalence scales obtained when allowing for heterogeneity with those obtained when imposing the restriction $\eta_l = \eta^*$ or $\gamma_{0l} = \gamma_0^*$ indicates that almost all of the difference in the equivalence scales is attributed to the latter restriction on the fixed portion of the group’s subsistence level. When compensating across demographic groups, unobserved heterogeneity in the marginal utility of consumption is largely averaged out and therefore results in negligible distortion. In contrast, decisions regarding the extent to which welfare equalisation accounts for differences in household-specific subsistence levels have a significant impact on equivalence scales.

6 The importance of household-level differences

The calculations in the preceding section indicate that differences in subsistence levels, rather than marginal utilities of consumption, are responsible for most of the demographic *group-level* distortion resulting from the assumption of homogeneity. This has important consequences for the design of effective income support policies. At first glance, the group-level results appear to provide evidence in favour of the welfare importance of heterogeneity driven by adjustment costs (Chetty and Szeidl, 2016), and against the importance of habit-driven heterogeneity (see, for example, Calvet and Comon, 2003). However, measurement of the unobserved characteristics that Pollak and Wales (1979) and Fisher (1987) highlight as relevant to any welfare comparison also requires an assessment of *intra-group* differences between households. It is these differences that provide information on what van Praag (1980) calls the ‘contour’ of income requirements that stem from household heterogeneity.

To see why intra-group unobserved differences are important, consider the scenario where the policy-maker chooses to ignore unobserved heterogeneity. In this scenario, the

equivalence scale for group l is determined by the appropriate transformation of (22)

$$\frac{\bar{y}_l}{y^*} | (\eta_l = \eta^*, \gamma_{0l} = \gamma_0^*) = 1 + \frac{\tilde{c}_l - \tilde{c}^*}{y^*} \quad (34)$$

where $\tilde{c}_l = x'_l \gamma$ characterises the observed heterogeneity of group l (with \tilde{c}^* characterising the observed heterogeneity of the benchmark group). It is clear that all households belonging to group l will have the same equivalence scale. This is the approach that is generally adopted.³

Conversely, the policy-maker who considers unobserved heterogeneity adopts the equivalence scale

$$\frac{\bar{y}_{il}}{y^*} = \frac{\eta_{il}}{\eta^*} + \frac{\hat{k}_{il} - \frac{\eta_{il}}{\eta^*} \hat{k}^*}{y^*} \quad (35)$$

where $\hat{k}_{il} = \gamma_{0il} + \tilde{c}_l$ and η_{il}, γ_{0il} pertain to individual households that belong to group l (with \hat{k}^*, η^* reflecting the analogous parameters for the benchmark 4-member household). In this case, the equivalence scale is able to account for heterogeneity in the marginal utility of consumption (through $\frac{\eta_{il}}{\eta^*}$) and in subsistence levels (through $\hat{k}_{il} - \frac{\eta_{il}}{\eta^*} \hat{k}^*$). If there is no such heterogeneity then the difference between $\frac{\bar{y}_{il}}{y^*}$ and $\frac{\bar{y}_l}{y^*} | (\eta_l = \eta^*, \gamma_{0l} = \gamma_0^*)$ will be zero.

Intra-group variation in η_{il}, γ_{0il} will produce a contour of equivalence scales for group l that is used to measure the welfare distortion of unobserved differences. To evaluate these distortions, it is instructive to consider the general appropriateness of grouping households by reference to family size as this grouping is ubiquitously adopted for the purpose of reporting equivalence scales (see, for example: Jorgenson and Slesnik, 1984 and 1987; Phipps and Garner, 1994).

Since households of a given size will differ in terms of observed features x , \tilde{c}_{il} is also observed for each household in group l and represents variation in observed heterogeneity. To restrict dispersion to that attributable to *unobserved* heterogeneity, a common value \tilde{c}_l is required across all households in group l . Accordingly, without loss of generality, \tilde{c}_l is set equal to \tilde{c}^* .

³If group l is characterised by a subset of x_l , then equivalence scales will only differ by reference to observed differences in the compliment of that subset with respect to x_l .

Figure 3 shows the entire distribution or contour of welfare distortions that are due to unobserved heterogeneity.⁴ Since we restrict \tilde{c}_i to be equal to \tilde{c}^* , the equivalence scale (34) is, by definition, always equal to unity such that the distributions in Figure 3 are also the distributions of the heterogeneous equivalence scale (35) after subtracting unity. In other words, with the exception of the mean, the distribution of the welfare distortions and the distribution of the heterogeneous equivalence scales are identical if we restrict our attention to unobserved differences.

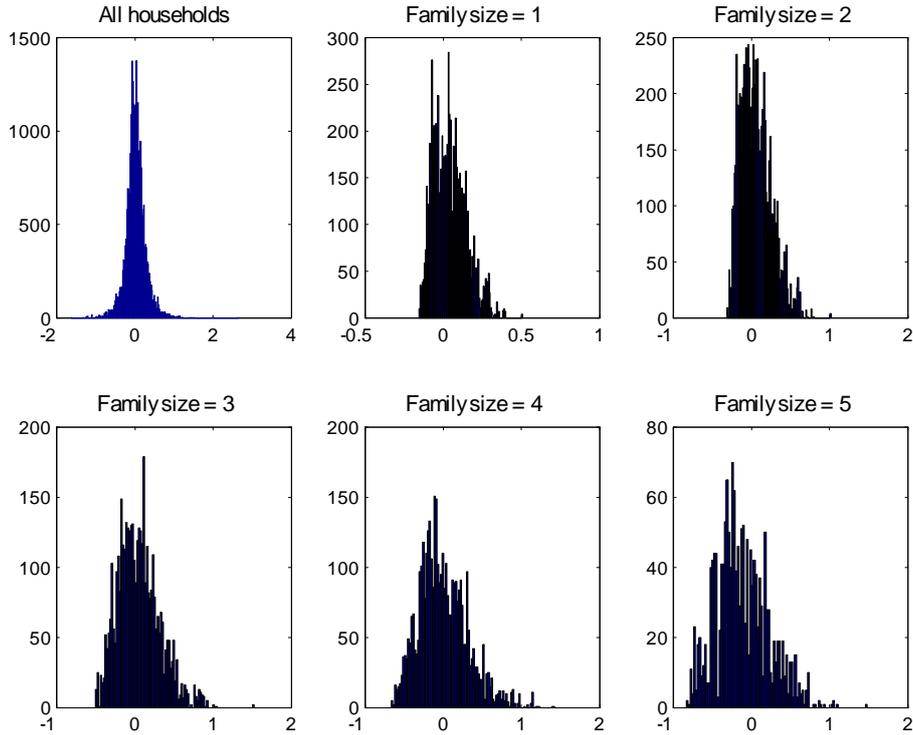


Figure 3: Distribution of welfare distortions D_l^* (for all households and by groups based on household size) stemming from unobserved heterogeneity. Positive values imply that the relevant households in the given group require additional income in order to attain the well-being level of the benchmark 4-person household.

For each family size, the level of distortion due to unobserved heterogeneity is substantial. Although differences in both marginal utilities and subsistence levels contribute to the distortion, the latter is responsible for the majority of the dispersion. The distributions clearly differ, however, across the various family sizes, suggesting that the welfare

⁴These are calculated as the difference between (35) and (34)

distortion associated with the failure to account for unobserved heterogeneity will be greater for some family sizes relative to others.

This result has significance for welfare analysis. It is a standard, and often implicit, assumption in welfare analysis that the ‘distribution of unconditional preferences is independent of the distribution of demographic characteristics’ (Pollak and Wales, 1979). The distributional differences across family sizes in Figure 3 render it relatively clear that this assumption is untenable. Indeed, the null hypothesis that the distributions in Figure 3 are from the same underlying distribution is rejected at the .01 level for *every* family size.⁵

The mean, variance and skewness of the welfare distortions

Table 4 specifies the centrality, variance and skewness of the distribution of heterogeneous equivalence scales. The variance and skewness have direct welfare interpretations if it is assumed that the policy-maker equalises welfare by adopting either a homogeneous equivalence scale (34) or the sample mean of the heterogeneous equivalence scale. In particular, the standard deviation represents the dispersion of welfare distortions stemming from unobserved heterogeneity, whereas the skewness provides information regarding the typical direction of the distortion. Furthermore, the mean absolute deviation represents the average welfare distortion. This can also be interpreted as a measure of the average *social welfare loss* when the policymaker assumes a transfer policy that does not account for household-level heterogeneity.

The impact of unobserved heterogeneity on the equivalence scales is substantial, and increases by family size. For families of size 3, 4 or 5 the average welfare distortion stemming from unobserved heterogeneity is approximately 25 per cent. This value rises to 35 per cent for families of size 6 and falls to about 8 per cent for a lone person. In all cases, significant positive skewness implies that the welfare distortions associated with the failure to account for unobserved heterogeneity are *asymmetric*, and typically result in under-compensation.

⁵The tests are based on the Kolmogorov Smirnov test and the resulting p-values are provided in Table 4.

Table 4: Summary statistics for the distribution of welfare distortions stemming from (i) unobserved heterogeneity and (ii) both observed and unobserved heterogeneity

	(i) Unobserved heterogeneity					(ii) Observed & Unobserved heterogeneity				
Family size	Mean	Std	MAD	Skew	Sig.	Mean	Std	MAD	Skew	Sig.
1	0.287	0.104	0.084	0.638	<0.01	0.285	0.120	0.098	0.491	<0.01
2	0.551	0.196	0.157	0.712	<0.01	0.547	0.225	0.182	0.547	<0.01
3	0.798	0.268	0.212	0.654	<0.01	0.806	0.294	0.231	0.522	<0.01
4	1.000	0.315	0.249	0.707	-	1.000	0.337	0.266	0.594	-
5	1.126	0.336	0.268	0.519	<0.01	1.121	0.377	0.300	0.422	<0.01
6	1.325	0.455	0.348	0.920	<0.01	1.329	0.493	0.375	0.798	<0.01

Note: The summary statistics for unobserved heterogeneity omit the affects of observed heterogeneity, whereas the summary statistics for ‘observed and unobserved heterogeneity’ account for both forms of heterogeneity. MAD is the mean absolute deviation which can be shown to be the average welfare distortion or social welfare loss. Significance tests are based on a Kolmogorov-Smirnov test that the sample of equivalence scales for family size j is from the same underlying distribution as the distribution for 4-person households.

The welfare importance of unobserved heterogeneity relative to observed heterogeneity

It follows naturally to consider the magnitude of unobserved heterogeneity relative to observed heterogeneity. If the impact of unobserved heterogeneity is small once observed heterogeneity is also accounted for, the results become substantially less important. To assess the magnitude of observed heterogeneity, the equivalence scales are re-computed after allowing \tilde{c}_i to vary across households (ie. by using the actual \tilde{c}_{il} for each household). For conciseness, the results are also presented in Table 4.

The change in the average distortion after allowing for observed heterogeneity (measured by the difference in mean absolute deviations across the two relevant columns) is relatively minor. In this respect, since the variance of the observed component of heterogeneity \tilde{c}_{il} must be greater than or equal to zero, it follows that the standard deviation of the equivalence scales when allowing for observed and unobserved heterogeneity will

be greater than or equal to the standard deviation obtained when allowing only for unobserved heterogeneity, hence $\sigma_{total} \geq \sigma_{unobserved}$. The proportion of the variability in the welfare distortions that can be attributed to unobserved heterogeneity can therefore be determined as the ratio of the standard deviations

$$\frac{\sigma_{unobserved}}{\sigma_{total}} \tag{36}$$

where $\sigma_{unobserved}$, σ_{total} are obtained from Table 4.

Pursuant to (36), just under 90 per cent of the heterogeneity in the equivalence scale is attributable to unobserved characteristics.⁶ The values exhibit relatively little variation across household sizes, ranging from a low of 86.7 per cent for single person households to a high of 93.4 per cent for four person households. The distortion stemming from unobserved heterogeneity therefore appears to be substantial, both in its own right and relative to the impact of observed heterogeneity.

Similar results are obtained even if the panel's inclusion criteria is modified from households that have been interviewed at least 6 times to households that have been interviewed 8 or 10 times. The observed result is also maintained when initially restricting the household lifespan to a minimum of 6 periods for all households and then repeating the estimation at each time point by adding an extra observation (if available).

Irrespective of the precision of the estimate provided by (36), however, the magnitude of the heterogeneity attributed to unobserved factors renders it clear that a large proportion of the heterogeneity in the equivalence scales is due to the unobserved aspects of each household, and that the failure to account for this heterogeneity is likely to produce significant welfare distortions (primarily in the form of under-compensation).

⁶The weighted average percentage of heterogeneity attributed to unobserved characteristics is 88.8 per cent.

7 Conclusion

It is shown that the equivalent income derived from popular models of well-being (such as probit or generalised extreme value models) can be mapped to the equivalent income stemming from utility functions that allow for household specific subsistence levels and marginal utilities of consumption. The paper then derives measures of the welfare distortion attributable to unobserved heterogeneity which can be computed using standard well-being models. These serve as a basis for examining the key issue raised in Pollak and Wales (1979) and Fisher (1987) regarding the distortions stemming from a failure to account for unobserved heterogeneity in household expenditure requirements.

The welfare distortion measures are applied to an Australian dataset which is used as a case study for examining the distribution of the welfare distortions. The centrality, variability and skewness of the resulting distributions are reported in order to determine the welfare implications of unobserved heterogeneity.

The results have important ramifications for understanding the impact of income and transfer policies, and the extent to which such policies induce distortions in welfare levels. In particular, the evidence broadly indicates that unobserved heterogeneity produces welfare distortions in two distinct ways.

First, a policy-maker who sets income requirements for particular household types without accounting for unobserved heterogeneity is likely to under-compensate smaller households, and over-compensate larger households. This is because the greater marginal utility of consumption of larger households is more than offset by their lower per-capita subsistence levels.

Second, the failure to account for unobserved household-level heterogeneity results in substantial and asymmetric welfare distortions, with the asymmetry typically resulting in the under-compensation of households. This is primarily due to the failure to capture differences in subsistence levels, rather than heterogeneity in marginal utilities of consumption.

References

- [1] Aguiar, M. and Hurst, E. (2005), "Consumption versus Expenditure," *Journal of Political Economy*, 113(5), 919-948.
- [2] Aiyagari, S.R. (1994), "Uninsured idiosyncratic risk and aggregate saving," *Quarterly Journal of Economics*, 109(3), 659-684.
- [3] Australian Bureau of Statistics (2004), "Australian Social Trends 2004," Cat. No. 4102.0, Canberra.
- [4] Australian Bureau of Statistics (2011), "Household Expenditure Survey 2009-10," Cat. No. 6530.0, Canberra.
- [5] Attanasio, O. and Weber, G. (1995), "Is consumption growth consistent with intertemporal optimization? Evidence from the Consumer Expenditure Survey," *Journal of Political Economy*, 103(6), 1121-1157.
- [6] Blackorby, C. and Donaldson, D. (1994), "Measuring the cost of children: A theoretical framework," in: R. Blundell, I. Preston, & I. Walker (eds.), *The measurement of household welfare*, Cambridge University Press, Cambridge..
- [7] Blundell, R., Pistaferri, L., and Preston, I. (2008), "Consumption inequality and partial insurance," *American Economic Review*, 98(5), 1887-1921.
- [8] Blundell, R., Kristensen, D., and Matzkin, R. (2014), "Bounding quantile demand functions using revealed preference inequalities," *Journal of Econometrics*, 117, 112-127.
- [9] Briesch, R.A, Chintagunta, P.K. and Matzkin, R.L. (2002), "Semiparametric Estimation of Choice Brand Behavior," *Journal of the American Statistical Association*, 97(460), 973-982.
- [10] Briesch, R.A, Chintagunta, P.K. and Matzkin, R.L. (2009), "Nonparametric Discrete Choice Models with Unobserved Heterogeneity," *Journal of Business and Economic Statistics*, 28(2), 291-307.
- [11] Browning, M., Chiappori P-A. and Lewbel, A. (2013), "Estimating Consumption Economies of Scale, Adult Equivalence Scales, and Household Bargaining Power,"

- Review of Economic Studies, 80(4), 1267–1303.
- [12] Browning, M and Crossley, T.F. (2001), “The Life-Cycle model of consumption and saving," *Journal of Economic Perspectives*, 15(3), 3-22.
- [13] Buera, F.J., Monge-Naranjo A., and Primiceri, G.E. (2011), “Learning the wealth of nations," *Econometrica*, 79(1), 1-45.
- [14] Calvet, L. and Comon, E. (2003), “Behavioural heterogeneity and the income effect," *Review of Economics and Statistics*, 85(3), 653–669.
- [15] Chetty, R. and Szeidl, A. (2007), “Consumption commitments and risk preferences," *Quarterly Journal of Economics*, 122(2), 831-877.
- [16] Chetty, R. and Szeidl, A. (2016), “Consumption commitments and habit formation," *Econometrica*, 84(2), 855–890.
- [17] Cherchye, L., I. Crawford, B. De Rock, and F. Vermeulen (2015), “Revealed Preference and Aggregation," mimeo.
- [18] Chiappori, P.A. (2016), “Equivalence versus indifference scales," 126(592), 523-545.
- [19] Clarida, R.H. (1991), “Aggregate stochastic Implications of the life cycle hypothesis," *Quarterly Journal of Economics*, 106(3), 851-867.
- [20] de Ree, J., Alessie, R., and Pradhan, M. (2013), “The price and utility dependence of utility scales: evidence from Indonesia," *Journal of Public Economics*, 97, 272-281.
- [21] Decancq, K., Fleurbaey, M. and Schokkaert, E. (2017), “Wellbeing Inequality and Preference Heterogeneity," *Economica*, 84(334), 210-238.
- [22] Donaldson, D. and Pendakur, K. (2003), “Equivalent-expenditure functions and expenditure-dependent equivalence scales," *Journal of Public Economics*, 88, 175-208.
- [23] Donaldson, D. and Pendakur, K. (2006), “The identification of fixed costs from consumer behavior," *Journal of Business and Economic Statistics*, 24, 255-265.
- [24] Dynan, K.E., Skinner, J., and Zeldes, S.P. (2004), “Do the rich save more?" *Journal of Political Economy*, 112, 397-444.
- [25] Engel, E. (1895), “Die Lebenskosten Belgischer Arbeiter-Familien Fruher and jetzt," *International Statistical Institute Bulletin*, 9, 1-74.

- [26] Fernandez-Villaverde, J. and Kruger, D. (2007), "Consumption over the life cycle: Some facts from consumer expenditure survey data," *Review of Economics and Statistics*, 89(3), 552-565.
- [27] Fisher, F.M. (1987), "Household equivalence scales and interpersonal comparisons," *Review of Economic Studies*, 54(3), 519-524.
- [28] Fisher, F.M. (1990), "Household equivalence scales: reply," *Review of Economic Studies*, 57(2), 329-330.
- [29] Goedhart, T., Halberstadt, V., Kapteyn, A., and van Praag. B.M. (1977), "The poverty line: concept and measurement," *Journal of Human Resources*, 12(4), 503-520.
- [30] Gorman, W. (1953), "Community preference fields", *Econometrica*, 21, 63-80.
- [31] Gorman, W. (1961), "On a class of preference fields", *Metroeconomica*, 13, 53-56.
- [32] Gourinchas, P.E. and Parker, J. (2002), "Consumption over the life cycle," *Econometrica*, 70(1), 47-89.
- [33] Jorgenson, D.W. and Slesnick, D.T. (1984), "Aggregate consumer behaviour and the measurement of inequality," *Review of Economic Studies*, 51(3), 369-392.
- [34] Jorgenson, D.W. and Slesnick, D.T. (1987), "Aggregate consumer behaviour and household equivalence scales," *Journal of Business and Economic Statistics*, 5(2), 219-232.
- [35] Koulovatianos, C., Schroder, C., and Schmidt, U. (2015), "Do demographics prevent consumption aggregates from reflecting micro-level preferences?" Working Paper.
- [36] Krueger, D. and Perri, F. (2006), "Does income inequality lead to consumption inequality? Evidence and theory," *Review of Economic Studies*, 73(1), 163-193.
- [37] Lusardi, A., Schneider, D., and Tufano, P. (2011), "Financially fragile households: evidence and implications," *Brookings Papers on Economic Activity*, 42(1), 83-150.
- [38] Lewbel, A. and Pendakur, K. (2017), "Unobserved Preference Heterogeneity in Demand Using Generalized Random Coefficients," *Journal of Political Economy*, 125(4), 1100-1148.

- [39] Mazzocco, M. (2007), "Household intertemporal behaviour: A collective characterization and a test of commitment," *Review of Economic Studies*, 74, 857-895.
- [40] Morciano, M., Hancock, R., and Pudney, S. (2015), "Disability Costs and Equivalence Scales in the Older Population in Great Britain," *Review of Income and Wealth*, 61(3), 949–514.
- [41] Phipps, S. and Garner, T.I. (1994), "Are equivalence scales the same for the United States and Canada?" *Review of Income and Wealth*, 40(1), 1-17.
- [42] Pollak, R.A. (1991), "Welfare comparisons and situation comparisons," *Journal of Econometrics*, 50(1-2), 31-48.
- [43] Pollak, R.A. and Wales, T.J. (1979), "Welfare comparisons and equivalence scales," *American Economic Review*, 69(2), 216-221.
- [44] Pollak, R.A. and Wales, T.J. (1981), "Demographic variables in demand analysis," *Econometrica*, 49, 1533-1551.
- [45] Senik, C. (2004), "When information dominates comparison: Learning from Russian subjective panel data," *Journal of Public Economics*, 88(9-10), 2099–2123.
- [46] van Praag, B.M. (1968), *Individual Welfare Functions and Consumer Behavior*, North Holland Publishing Co., Amsterdam.
- [47] van Praag, B.M. (1971), "The welfare function of income in Belgium: an empirical investigation," *European Economic Review*, 2, 337-69.
- [48] van Praag, B.M., Goedhart, T., and Kapteyn, A. (1980), "The poverty line - a pilot survey in Europe," *Review of Economics and Statistics*, 62(3), 461-465.
- [49] Zaidi, A. and Burchardt, T. (2005), "Comparing Incomes When Needs Differ: Equivalization for the Extra Costs of Disability in the U.K.," *Review of Income and Wealth*, 51(1), 89–114.

Appendix: Covariates

Table A: Covariates x_{it} used to estimate the income requirement k_{it}

employee	promoted	serious injury to family member
self-employed	back together with spouse	serious personal injury
employer	retired	family member jailed
unpaid worker	year = 2003	jailed
# jobs in last FY	year = 2004	victim of property crime
long-term health condition	⋮	separated
renter	year = 2012	victim of physical crime
mortgagee	death of close friend	
changed jobs	weather disaster	
married	death of close relative	
changed residence	death of spouse	
pregnancy	fired from job	

