IZA Institute of Labor Economics

Initiated by Deutsche Post Foundation

DISCUSSION PAPER SERIES

IZA DP No. 13326

Inequality Measurement and Tax/Transfer Policy

Patricia Apps
Ray Rees

JUNE 2020

I Z A Institute of Labor Economics

Initiated by Deutsche Post Foundation

DISCUSSION PAPER SERIES

IZA DP No. 13326

Inequality Measurement and Tax/Transfer Policy

Patricia Apps
University of Sydney Law School and IZA

Ray Rees
University of Munich, University of Sydney Law School and CESifo

JUNE 2020

ISSN: 2365-9793

IZA – Institute of Labor Economics

Schaumburg-Lippe-Straße 5–9
53113 Bonn, Germany

Phone: +49-228-3894-0
Email: publications@iza.org

www.iza.org

ABSTRACT

Inequality Measurement and Tax/Transfer Policy

We provide a critique of the standard methodology which bases welfare comparisons between households on deflating household income and consumption by an equivalence scale. We argue that this leads to support for tax/transfer policies that significantly disadvantage low to middle income households and women as second earners. We base the critique both on a theoretical model of the family household and a detailed analysis of Australian income and employment data.

JEL Classification: D13, D31, H21, H24, H31
Keywords: inequality, equivalence scales, tax/transfer policy, families

Corresponding author:
Patricia Apps
The University of Sydney Law School
The University of Sydney
NSW 2006
Australia
E-mail: patricia.apps@sydney.edu.au

1 Introduction

The use of equivalisation indices to deflate measures of household income and consumption in empirical studies of the family household has become widespread and routine, although some economists have argued strongly against it.[1] In practice a number of widely used equivalence scales exist. Their aim is based on an assumed need to take account of the variations in size and composition of households and the economies of scale in household production when making welfare comparisons across heterogeneous households.

For example the widely used OECD "square root" scale[2] deflates household aggregates, such as gross and disposable incomes and total consumption, by the square root of the number of individuals in the household. These are then used to construct indices of inequality across entire economies. Another typical example is the "Oxford modified" scale used by the Australian Productivity Commission.[3] A scale of 1 point is used for the first adult, 0.5 for each additional person aged 15 years or more, and 0.3 points for each child under 15 years. A number of other scales that use a similar procedure, but with different numbers, have been proposed over the years. The scales are constructed on the basis of some calculation of the "needs" of individuals of different ages. Then with each household in the sample is associated a number of "adult equivalents". The idea of economies of scale in household production, often expressed in the old adage "two can live as cheaply as one", is reflected in weights of less than one for adults beyond the first. Thus a family of two adults with two children is considered under the Oxford modified scale to be equivalent to 2.1 single adults. By deflating a household's total income by the equivalence index and assuming that income is equally distributed across adult equivalents within the household, a single number is obtained, which, it is assumed, characterises the standard of living of each adult in the household and is directly comparable across households. If, for example, a household with two adults and two children had a joint income of 2.1 times that of a reference single person household the two households would be considered equally well off. We do not find it difficult to be critical of this type of procedure as a way of dealing with household heterogeneity, as we shall make clear in the remainder of this paper.

The simple assumption of "economies of scale" applying in some way to household well-being does not do justice to the complexity of realistic household production processes - see for example the critique of the Becker model of household production in Pollak and Wachter (1975), who emphasise the importance of multi-activity production functions characterised by significant joint production,[4] without the presumption that economies of scale apply to all forms

[1] For example see Pollak and Wales (1979) who completely rejected the procedure, and Atkinson (1970) who criticised the "needs based" approach to their construction.

[2] See, for example, OECD (2019) and Sila and Dugian (2019).

[3] Report of the Australian Productivity Commission (2018).

[4] The fact for example that one can be looking after a small child as well as carrying out other household tasks is often ignored in time use studies. Most report only one time use for each activity episode. An exception is the Australian 2005-06 Time Use Survey which reports two activities for each activity episode, but has not been continued. More recent US surveys

of household production.[5] In this paper however our main focus will be on the use of total household income as the base for the equivalisation process, and of the *ad hoc* assumption, completely unsupported by empirical evidence, of complete equality in the within household distribution of income and consumption. This is the result of the absence of any attempt to model the resource allocation decisions made within households, and the explicit model of these we present in this paper[6] suggests that the assumption is not in general tenable.

In the more theoretical literature attempts have been made to put these procedures on a less *ad hoc* basis.[7] Angus Deaton and John Muellbauer,[8] at the end of a brilliantly clear exposition of the theoretical fundamentals of the approach (which they basically favour on the grounds of its econometric advantages), write:[9]

> We note, finally, a source of some difficulty in the treatment of children [in the economic analysis of the household]. *So far we have cavalierly ignored the distinction between households and individuals, treating the two terms more or less interchangeably.** Our preference in analysing behaviour is to treat the *household*** as the basic decision-making unit, modeling the behavioral impact of family composition through the equivalence scale [...]. But this is not entirely satisfactory. Social welfare is formed over individual welfares so that society is not likely to be unconcerned about how members of families are treated [...]. The social welfare function should thus have a "slot" for each individual, and *if each family member has the same welfare level,** the family per capita equivalent real income can be used as the welfare indicator for each individual. [This approach is] *not entirely satisfactory without a theory of (or at least some assumption about) allocation within the household.**
>
> *Our italics. **Italics in the original

We should first point out that the authors are ignoring not only the "treatment of children" but also gender differences in couple households, which of course are closely related to the presence of the children. This is reinforced by the absence of any consideration of household production. Moreover, it would surely be missing the point to base an analysis involving the equity of the distribution of income and consumption across *individuals* on an assumption that the household's sharing rule involves equality for all its members.

report if a child is present which allows, to some extent, misleadingly low estimates of the time a household spends on child care to be avoided.

[5] Parents with two or more small children would certainly doubt that child care exhibits increasing returns to scale.

[6] This is of course part of a large body of relatively new literature which takes the family household seriously. See for example Gayle and Shephard (2019), Lundberg et al. (2016), Alesina et al. (2011), Konrad and Lommerud (2000), Lundberg et al. (1997), and Basu (2006). For more comprehensive literature surveys see Apps and Rees (2009) and Bergstrom (1993).

[7] See Appendix 1 for an outline of the Deaton/Muellbauer analysis.

[8] Deaton and Muellbauer (1980).

[9] Deaton and Muellbauer (1980), p 226.

Our aim in this paper is to show that household equivalence scales are both unnecessary and misleading. They give support to tax/transfer policies that make low-to-middle income households, and women in those households as second earners, significantly worse off than they would be under alternative policies. As a case in point, we take the withdrawal of child support payments on the basis of a household's joint income, which transforms personal income taxation into a quasi-joint rather than an individual income-based system.[10] Following the presentation of the model giving the conceptual framework for the discussion, we present an in-depth analysis of the Australian family tax system to illustrate our argument.

2 A critique of equivalisation indices

This critique is based on a model of household decision-taking behaviour.[11] We use this model[12] to argue that, as mentioned in the Introduction, there are some fundamental weaknesses of the equivalisation index approach that undermine its usefulness in applications to policy.

First, the procedure of using indices to deflate household joint income embodies the idea that household income and well-being are co-monotonic - one necessarily increases with the other as we move through the equilibria of the given set of households.[13] In fact, given the marked heterogeneity in second earner labour supply, which has been firmly established empirically, joint labour market income is an inaccurate and misleading indicator of household wellbeing, as we show below. The root of this misconception lies in ignoring the existence of household production in a multi-person household.

Secondly, the equivalisation procedure, when based on standard demographic variables, assumes that the components of the household type vector[14] are fully observable, and therefore rules out consideration of the implications of the fact that important components of this vector are not observable, or at least not available in existing datasets. This has to do not only with the omission of

[10] For a powerful critique of targeting and a thorough investigation of better alternatives see Atkinson et al. (2017).

[11] We draw here on the model of the household as a small economy engaged in intra-household production and exchange in Apps (1982). In the original formulation this was a general equilibrium model in which market wage rates were endogenous. The gender wage gap was driven by the "crowding" of women into "female" occupations, which leads to a higher male wage and lower implicit price for the household good. This drew on the model of racial discrimination in Bergmann (1971). Here we simplify by assuming that the terms of this exchange are determined exogenously.

[12] Set out fully in Appendix 2.

[13] This is true not only of the sophisticated measures such as that in Deaton and Muellbauer, but also of the needs-based counting measures, such as the Modified Oxford procedure, which compute the index by attaching a number to each individual and then adding them up.

[14] This is the list of exogenously given household characteristics that can cause two households with the same total incomes and preferences to choose different consumption bundles and time allocations. Important components of these vectors are wage rates. They may also include some prices, for example, of bought in goods used in household production.

household production, but more fundamentally with the absence of a conceptual framework for household decision taking of the kind we provide in this paper.

2.1 The household model

In the model, a household consists of two adults, each of whom works in the labour market and also supplies time to production of a household good, which could, but need not, be thought of as child care.[15] They also consume some of their own time as leisure. They have standard individual utility functions

$$u_{ih} = u(x_{ih}, z_{ih}, l_{ih}) \quad i = 1, 2 \qquad h = 1, 2, ..., H \tag{1}$$

where 1 is the primary and 2 the second earner.[16] Consumptions of a composite market good, the domestic good and leisure[17] are respectively x, z and l, all strictly positive. By definition, 1 has the higher labour market earnings. Note that we assume identical preferences within and across households.

The household production function, with $\sum_{i=1}^{2} z_{ih} = z_h$, is

$$z_h = f(k_{1h}a_{1h}, k_{2h}a_{2h}, q_h b_h) \qquad h = 1, 2, ..., H \tag{2}$$

where a_{ih} is i's time input into household production and b_h is a market input bought at price p_h. Note that this price may vary across households.[18] Moreover, the k_{ih}, q_h are productivities/qualities of the inputs into household production that may also vary across households.[19]

The individual time constraints are

$$a_{ih} + l_{ih} + L_{ih} = T \quad i = 1, 2 \tag{3}$$

where T is total time available and $L_{ih} \geq 0$ are market labour supplies.

A key issue in formulating a household model is the budget constraint. Necessary (though not sufficient) conditions for the assumption that in any household equilibrium utilities of its members are equalised are two conditions:

[15] In fact the model here is not fully applicable to households in the life cycle phase in which children under school age are present, since in that case additional constraints are placed on time use - someone always has to look after the kids. For further discussion and analysis see Apps and Rees (2018).

[16] Note that these are defined on role rather than gender, but empirically in OECD countries typically around 80% of second earners are women.

[17] For convenience of notation we assume that these 3 consumption goods are scalars rather than vectors, which in turn implies that, since x is the numeraire with price 1, all other prices and wage rates have been deflated by a consumption price index. This is without real loss of generality. A more important omission is that of corner solutions involving zero labour supply for second earners. We do however discuss this case less formally below.

[18] For example, if the input b_h were non-parental child care, there could be a wide range of sources, such as family members and friends, child-minders, childcare centres and posh nannies.

[19] For example, it is well-established in the literature that the quality of childcare - in the broad sense of creating human capital as well as simple physical child-minding - depends on parental human capital as well as other resources possessed by the household. See Apps and Rees (2018) for references.

1. transfer payments are feasible - utility is transferable
2. *only* the *pooled* household budget constraint must hold:

$$\sum_{i=1}^{2} x_{ih} + p_h b_h \leq \sum_{i=1}^{2} w_{ih} L_{ih} \quad h = 1, 2, ..., H \tag{4}$$

where w_{ih} are the market wage rates, with $w_{1h} \geq w_{2h}$.[20]

In this paper we want to move away from the case in which individual welfares are equalised and so we make the alternative extreme assumption that within-household lump sum transfers are ruled out.[21] We assume that individuals maximise utility subject to budget constraints determined by their own full incomes, given by the value of their time endowments at their own market wage rates, $w_{ih}T$. This gives us a simple and observable sharing rule that determines the within-household distribution of utility. The values of the individual total consumptions of the market and the household goods as well as leisure are determined by their full income. We then have the basis for the following analysis of inequality *across individuals* and its measurement.

The household is fully rational in that it values each earner's time consistently at their outside market wage. In particular it prices individual leisures at their corresponding market wage and applies the implicit price π_h as the marginal opportunity cost of each individual's consumption of the household good.[22] This implies that we can represent the individual choice problems in terms of their *full income budget constraints*:

$$\max_{x_{ih}, z_{ih}, l_{ih}} u(x_{ih}, z_{ih}, l_{ih}) \text{ s.t. } x_{ih} + \pi_h z_{ih} + w_{ih} l_{ih} \leq w_{ih} T \quad i = 1, 2, \quad h = 1, 2, ..., H \tag{5}$$

This emphasises that the wage rates w_{ih} and the prices π_h jointly determine the utility possibilities of all individuals in all households. From this we could define the type vector in this model as $[w_{1h}, w_{2h}, \pi_h]$ since these are the exogenous variables that determine the household equilibrium. However, the implicit price of the household good is itself a function given by[23]

$$\pi_h = c(w_{1h}, w_{2h}, p_h, k_{1h}, k_{2h}, q_h) \tag{6}$$

and so we would define the household's type vector as $[w_{1h}, w_{2h}, p_h, k_{1h}, k_{2h}, q_h]$.

[20] The underlying assumption here is that the household chooses its role specialisation on the basis of wage rates. This is consistent both with the gender wage gap and with the fact that the majority of second earners are women.

[21] This assumption is essentially simplifying. If we are to allow within household transfers, we have to assume the existence of a household welfare function that gives the household's preferences over utility levels of its members, for example a Nash bargaining function, or more generally a weighted utilitarian function as in Apps and Rees (1988) and Chiappori (1988), now usually referred to as the "collective model".

[22] For details of the derivation of this price see Appendix 2.

[23] See Appendix 2.

We now consider the indirect utility and expenditure functions that result from this model, and how they compare to those delivered by the model of Deaton and Muellbauer. We can write them as, respectively:

$$v_{ih} = v(w_{1h}, w_{2h}, p_h, k_{1h}, k_{2h}, q_h) \tag{7}$$

$$e_{ih} = e(w_{1h}, w_{2h}, p_h, k_{1h}, k_{2h}, q_h) \tag{8}$$

First, as already noted, the relevant functions belong to individuals rather than households. This is then consistent with an approach that says social welfare functions should be defined on individuals rather than on collectives such as households. This is important in that it *removes the need for the ad hoc kinds of calculations that try to adjust for differences in household size and composition across an entire population.* Utilities are already individualised. The number and ages of children are taken into account in the variable z. Underlying this is the idea that children are a source of utility in the household, both to themselves and, one hopes, to their parents.[24] The issue of "needs" created by children is dealt with in the price π_h, which will be higher when children are present. This raises the issue of life cycle phases, taken up more fully in the next subsection.

Second, it removes household income, an endogenous variable, from the index measure, since the relevant functions are defined entirely on exogenous variables.

Third, it takes into account household production and the very significant heterogeneity of second earner labour supplies that results from that. As we see from the analysis in Appendix 2, labour supplies are given by the functions

$$L_{ih} = L_i(w_{1h}, w_{2h}, p_h, k_{1h}, k_{2h}, q_h) \quad i = 1, 2, \quad h = 1, 2, ..., H \tag{9}$$

and so the model explains this heterogeneity simultaneously and consistently with its explanation of utility variation.

Therefore we would propose on theoretical grounds either one of the indirect utility function $v_{ih}(.)$ or, if a money metric is preferred, the expenditure function $e_{ih}(.)$ as the measure of the well-being or standard of living of the individuals within the household, and the basis of a welfare ranking across individuals and households. These functions have been extensively used, though not in this form, in the analysis of optimal tax/transfer policies in public economics.

For the empirical work in this paper however the problem we face is that the data needed to implement this proposal does not yet exist, and so until it does some alternative has to be found. It seems to us that for employed couple households primary earner income is a readily available and conceptually suitable, if second best, substitute. Given that across this subset of households the data indicate that primary earner labour supplies do not vary widely, we take this as a good proxy for the primary earner's wage rate in the empirical analysis

[24]More to the point, children have their own utilities and public policy, for reasons related to market failure, is at least as much concerned with these as with the wellbeing of their parents. An alternative approach could specify children as individuals with their own utility functions in the model, but since household decisions are typically taken by adults with the perceived utility of children as arguments of their own utility functions the approach given here may be an acceptable reduced form.

to follow. With a significant and increasing degree of assortative matching it is also a reasonable proxy for the second earner's wage rate, particularly where second earner's incomes are relatively low and therefore dividing them by hours worked to estimate wage rates is likely to be subject to high measurement error.[25]

As we have stressed, what matters to household well-being is not simply second earner income, but that plus the value of her domestic output. The implicit assumption that the latter is zero, which underlies the use of household income as a measure of well-being, is, as we have argued, simply too extreme. We prefer to make the assumption that the sum of second earner income and value of household production is strongly positively correlated with primary earner income, which therefore makes primary earner income an even better proxy for household well-being.

Finally, given the lack of data on the type variables k_{1h}, k_{2h}, q_h, we would expect that the physical and human capital the household possesses is strongly positively correlated with the quality and productivity of the inputs into household production, and therefore can also be assumed to increase with primary earner income. Empirically, the availability and cost of bought in inputs into household production, once quality is controlled for, are probably best thought of as being randomly distributed across households.

Before proceeding with the empirical analysis based on these assumptions it is important to discuss issues concerning the life cycle of family households.

2.2 The importance of the life cycle

Welfare comparisons based on household income across a population of households that are at different phases of their life cycles are of very limited usefulness: what does it mean to compare on the basis of equivalised incomes the well-being of a young couple who have not yet had children, a couple with two children under school age, an older couple, both working, whose children have left home, and a retired couple? Of course, public policies require trade-offs to be made between the well-being of different subsets of the population, but the point of our critique is that this should not be done in the form of deflation of incomes or consumption by an equivalisation index. Essentially, the use of these procedures just sweeps the key issues - the welfare comparison of different household types for purposes of public policy - under the carpet. It is also often argued that the ideal would be the comparison of households' well-being over their entire life cycles, but this seems to be an impossible ideal to achieve, especially for policy purposes.

[25] It might also be argued that wage rates themselves are unobservable.Two points can be made in response to that. In tax analysis of the Mirrlees (1971) type "wage rates" are said to be non-observable, but this is a semantic issue. What is meant there is the non-observability of the innate ability of a worker, which in perfect labour markets is reflected in the wage. In this paper we are concerned with actual wage rates, however determined. A problem does exist however in measuring the opportunity cost of time for non-participating second earners, since any estimated market wage based on human capital variables gives only a lower bound on their marginal value product in household production.

When the principal aim of welfare comparisons is to design tax/transfer policies a different approach is required. Quite obviously, there has to be differentiation between the policy issues around households at different phases of the family life cycle. We would go further and argue that households should be assigned to life cycle phases according to the similarity of the preference structures and constraint sets that confront them, since these are the policy-relevant determinants of their well-being. For this reason, we have argued for a specification of the life cycle based not on the age of the "head of the household",[26] as is the usual case, but rather upon the phases through which family households typically go.[27]

For couples, for example, perhaps the broadest characterisation would be of a life cycle defined on 5 phases: the first phase in which the couple household has been formed but no children are present; the second, in which there are children of pre-school age in the household; the third, in which the children are older but still present in the household; the fourth, in which children have left home but the parents are still of working age: and the last phase in which both adults have retired.[28] Clearly, across these phases preferences and constraints can be expected to vary significantly, as would the vectors of characteristics exemplified by $(w_{1h}, w_{2h}, p_h, k_{1h}, k_{2h}, q_h)$ in the above analysis. Across families within each phase inequality across individuals and households based on our $v_{ih}(.)$ or $e_{ih}(.)$ measures should be applied.

It is also important to recognise that there are other household types that do not fit into our characterisation of the family life cycle, such as single-parent households, households with members suffering from disabilities, and single-person households. Since these differ in terms of observable characteristics they can be "tagged"[29] and made subject to our proposed methodology in a straightforward way.

Our argument is supported by the results for the properties of optimal piecewise linear income tax systems for families, based on joint and individual incomes, in Apps and Rees (2018). Drawing on data for families with a child under 10 years,[30] and therefore predominantly in phase 2 of the life cycle, the model shows that a small variation in the price of child care at a given wage pair can lead to wide variation in second earner labour supplies with little effect on welfare levels, when home and bought-in child care are close substitutes. A key implication of this finding is that primary income is likely to be a more reliable measure of household welfare than household income deflated by an equivalence scale because it implicitly takes account of the inverse relationship between the

[26] The discussion of "age-related taxation" though relevant for policy is misguided in taking calendar age rather than life cycle phase as the basis for categorising households.

[27] For a fuller discussion see Apps and Rees (2009), Ch 5. The approach suggested here could not be adopted if fertility were treated endogenously.

[28] See Apps and Rees (2011), Ch 3, for an analysis of time use and incomes based on a lifecycle model defined on these five phases.

[29] For further discussion of "tagged" payments, see Section 4.2.

[30] The sample is also limited to records with a primary earner aged 25 to 59 years and employed for at least 25 hours/week.

second earner's labour supply and her contribution to the family's standard of living through home production. This is a centrally important issue in tax design because of the high degree of heterogeneity in second earner labour supply at a given primary income,[31] as evidenced by the participation rates reported in the empirical analysis to follow. A further reason for its importance is that second earner labour supply decisions taken on the arrival of children in the household can have very significant, persistent effects on labour supply over the remainder of the life cycle.

In summary: On the basis of the results just discussed we argue that for the employed working age population, primary income is a far sounder basis for ranking households in terms of their well-being than equivalised joint income, even as a provisional second best while we await the collection of more data. Under the equivalised joint income measure, two households with the same number of children of similar ages, but in one of which two adults work full time to earn the same market income as that of the primary earner in the other, would still be regarded as equally well off. This ignores the fact that the second household will almost certainly be enjoying a much higher living standard because of its higher implicit income from household production and lower expenditure on its market substitutes. *Equivalisation does nothing to correct the fundamental weaknesses of joint income as the measure of a household's well-being.*

3 Empirical analysis: equivalised income

For the purpose of highlighting the limitations of equivalised income as a measure of household welfare, we present an analysis based on data for two-parent families with two dependent children. We select couple income unit records from the Australian Bureau of Statistics (ABS) 2015-16 Surveys of Income and Housing (SIH2015-16) on the following criteria: the primary earner is aged from 20 to 64 years, earns a minimum of $12/hour and works for at least 30 hours per week, and the couple income unit contains two dependent children both aged under 15 years. The sample contains 1055 records.

Section 3.1 begins by tabulating participation rates across quintile rankings defined first on primary income and then on equivalised income calculated according to the Oxford modified scale. Given that the sample is limited to couple income units with two children under 15, we have for each adult: equivalised income = household income/2.1. The results show a dramatic re-ranking of households when we switch from primary to equivalised income as the ranking variable. We investigate this ranking further in Section 3.2 by splitting the sample into two subsets defined according to the midpoint of second earner hours within each quintile.

[31] Important empirical papers on the relationship between taxation and family labour supply are Bick et al. (2018), LaLumia (2008) and Gayle and Shephard (2019).

3.1 Participation rates

Rows 1 and 2 of Table 1 report the quintile data means of primary and second incomes across a ranking defined on primary income. The next three rows contain second earner participation rates: non-participants (0 hrs), part-time (1-34 hrs) and full time (35+). Table 2 presents results for the same variables across a quintile ranking defined on equivalised income.

Tables 1 and 2 about here.

Overall, almost half of second earners, 45%, are working part-time, 27% are non-participants and 28% are in full-time work. There are, however, major differences in the quintile profiles of participation rates across the two tables. In Table 1 there is little variation in the distribution of participation rates across primary incomes. In contrast, when in Table 2 we take quintiles of equivalised household income, non-participants are heavily concentrated in the lower quintiles. This variation in the profiles of participation rates across quintiles, together with the gap in the data means for second incomes, reflects the high degree of re-ranking of households when we switch from primary to equivalised income as the welfare ranking variable.

The percentage of single-earner households in *quintile 1* rises from 29% in Table 1 to 54% in Table 2, and the proportion of full time second earners falls by more than half. In *quintile 5* the non-participation rate falls by more than half, and the proportion of full time second earners rises from 21% to 41%. The relatively flat quintile distribution of full-time second earners that we see in Table 1 becomes a sharply increasing distribution in Table 2.

Also evident in the comparison of the quintile distributions in the two tables is the compression in the distribution of primary incomes as we go from Table 1 to Table 2: the ratio of income in the top quintile to that in the bottom is nearly 20% larger in Table 1 than in Table 2, so that the re-ranking of *households* in Table 2 gives a more benign impression of the inequality in incomes across *individual* primary earners than is evident in Table 1. In fact the Gini coefficient for the primary income distribution in Table 1 is 0.330 and that for this distribution in Table 2 is 0.310.

3.2 Household subsets defined on second earner hours of work

To show more clearly the dramatic re-ranking of households when we switch from primary to equivalised income as the ranking variable and, in Section 4, the tax policy implications of that, the records in each quintile of primary income in Table 1 are split into two equal subsamples labelled H1 and H2:

H1 households: Second earner is working at or below the 50th percentile of the distribution of hours worked by second earners

H2 households: Second earner is working above the 50th percentile.

Next, we split records in each quintile of equivalised incomes in Table 2 into two subsamples, labelled E1 and E2, according to the same criterion:

E1 households: Second earner working at or below the 50th percentile

E2 households: Second earner working above the 50th percentile.

Table 3 reports H1 and H2 data means for primary incomes, second incomes and household incomes across quintiles of primary income. Table 4 presents E1 and E2 data means for the same income variables across quintiles of equivalised income.

Tables 3 and 4 about here

In Table 3, H1 and H2 *primary incomes* are strongly matching across quintiles 1 to 4, with a somewhat wider gap in quintile 5, where H1 primary incomes are significantly higher, an outcome associated with lower second earner participation in these households. H2 *second incomes*, while much lower than primary incomes, are around 4-5 times those of H1 households, as we would expect from the participation rate profiles in Table 1. H2 household incomes are of course consistently higher than those of H1 across the primary income distribution.

In contrast, in Table 4, E1 primary incomes are significantly higher than those of E2 within each quintile while of course E1 second incomes are markedly lower than those of E2. E1 and E2 household incomes are closely matching across the middle of the distributions, but there are more significant differences in the first and fifth quintiles.

The high degree of re-ranking associated with switching from the primary income ranking in Table 3 to the equivalised income ranking in Table 4 is evident from the distribution of H1 and H2 households according to *equivalised income* reported in Table 5. The percentage of H1 households in quintile 1 rises from 50% to 76% and in quintile 5, falls from 50% to 28%, while the percentage of H2 households in quintile 1 falls from 50% to 24% and rises from 50% to 72% in quintile 5. This outcome is shown graphically in Figure 1, where the sharply falling proportions of H1 households as we move across quintiles is reflected in the sharply increasing proportions of H2 households.

Table 5 and Figure 1 about here.

It is important to note that the strong tendency for the polarisation of H1 households towards the lower quintiles and of H2 towards the upper quintiles of equivalised income is driven not only by the high degree of heterogeneity in second earner labour supply, but also by the relatively flat profiles of primary incomes over the first 4 quintiles, with a very sharp increase in quintile 5, as indicated in Table 3. These in turn reflect the underlying wage profile graphed in Figure 2 below. When there is little variation in primary incomes below the top percentiles, a single income household in quintile 1 of equivalised income can be shifted towards quintile 4 when the second partner goes out to work for the same income. As a consequence, the household is regarded as being "upper middle income" and may be taxed accordingly, as the analysis to follow will illustrate.

4 Empirical analysis: Joint vs. individual taxation

The Australian family tax system provides an especially useful case study on which to base our analysis of individual vs. joint taxation, containing as it does elements of both. The aim of the analysis to follow is to show how equivalised income inequality measures boost support for joint taxation. The system combines the Australian individual based income tax with a family payment system, labelled Family Tax Benefit - Part A (FTB-A), which withdraws the payments on the basis of joint income.[32] The rate scale under the individual based income tax system is set by that of the Personal Income Tax (PIT) rate scale in combination with the Low Income Tax Offset (LITO). The PIT scale is strictly progressive but when combined with the LITO is no longer so, as shown below. FTB-A was introduced under the Howard Government's tax reforms of 2000 as part of a longer term agenda of shifting towards joint taxation with much less progressivity in the marginal tax rate scale on individual incomes,[33] together with the complete elimination of the earlier "universal family allowance" per child.[34]

Such reforms clearly change the effective tax parameters faced by each taxpayer, whether actual or potential. Under a piecewise linear muliple-bracket income tax system with varying marginal tax rates, each individual faces just two tax parameters: an effective marginal rate and an effective lump sum,[35] where the latter is the universal payment plus the positive (negative) "intra-marginal lump sum"[36] associated with lower (higher) MTRs across preceding income tax brackets. The effective marginal rate is the formal rate plus the rate at which transfer payments such as child benefits are withdrawn. The key point

[32] The system also includes Family Tax Benefit Part B which provides a payment that is withdrawn on the second earner's income. This element is omitted here in order to focus more clearly on the key implications of combining the Australian individual based income tax with FTB-A. The Medicare Levy is also omitted.

[33] Reforms prior to 2000 included large cuts in top marginal tax rates and post 2000, upward shifts in the bracket limits to which the top rate applied. As an example of the latter, from 2004-05 to 2008-09, the bracket limit at which the top marginal tax rate applied rose from $70,000 to $180,000.

[34] Initially, the size of the "universal family allowance" per child was reduced in real terms by the failure to index the payments over several decades. Their "universality" was completely eliminated under Howard Government reforms post 2000.

[35] The tax design problem is specified formally as that of solving for the optimal values of these tax parameters, together with the optimal bracket limits, as set out in Apps, Long and Rees (2014).

[36] The "intra-marginal lump sum" is calculated by subtracting the amount of tax actually paid from the amount that would be payable had the individual's total income been taxed at the marginal rate applying to her/his last dollar, and may be positive or negative depending on the progressivity/regressivity of preceding MTRs. The effective lump sum is then obtained by adding the result to any formal lump sum, such as family payments under FTB-A. In the simplest two-bracket system with formal lump sum a, marginal tax rates t_1, t_2, and bracket limit $\hat{y}$, the tax bill for lower rate taxpayers with income $y \leq \hat{y}$ is $t_1 y - a$ and for higher rate taxpayers it is $t_2 y - \hat{a}$, where $\hat{a} = a + (t_2 - t_1)\hat{y}$ is the effective lump sum.

is that "income targeting" a family payment or benefit, on either an individual or joint basis, does not end the payment of an effective lump sum. The immediate impact is a change in the marginal rate scale, typically making it not only less progressive but also no longer strictly progressive, if it raises the effective tax rate in lower brackets above the marginal tax rate in higher tax brackets in which the benefit has been totally withdrawn. Switching to income testing of the family payments on the basis of *joint* income in a previously *individual* income-based tax system changes the tax base of the system across the income withdrawal range of the benefits from individual to joint.

Support for targeting benefits typically draws on the specious argument that it is necessary to avoid making transfer payments to well-off households - so-called "middle class welfare" - and that it leads to a more efficient use of tax revenue. In the case of income targeting on the basis of joint income, the outcome cannot be an increase in economic efficiency. A higher tax rate on the second income cannot achieve efficiency gains given the evidence that second earner labour supply elasticities are significantly higher than those of prime age males in full employment, and there is no evidence that labour supply elasticities increase across the primary earner wage distribution. In the model set out in Section 2, higher second earner labour supply elasticities can be explained by a high elasticity of substitution between home production and market work, particularly in the case of home child care and bought in care.[37]

It could, of course, be argued that the efficiency loss associated with higher MTRs on second incomes can be avoided, or at least minimised, by basing the withdrawal rate on primary earner income. Interestingly, when the U.K. government introduced means-testing of child benefits, it used primary income as the basis to which the withdrawal rate is applied. In this case, while the tax system remains based on individual incomes, it can be classified as "selective": primary and second incomes are taxed under different rate scales. The primary earner faces a new rate scale, one with higher marginal rates across middle primary incomes. This case highlights what can be seen as a key driver of support for income testing benefits: that of replacing an existing progressive rate scale by one with higher marginal rates across middle incomes, combined with lower top rates, in other words, a shift to a much less progressive income tax system. This outcome raises the question as to whether targeting on the basis of either individual or joint incomes can be supported in terms of fairness. To answer this question we need to turn to the overall distribution of wage rates.

We draw on data for "in-work" couple income units in the ABS SIH2015-16 to construct the wage profiles in Figure 2.[38] The figure plots average primary and second earner wage rates across percentiles of primary income. The primary wage in each percentile is calculated as average gross earnings, with hours smoothed across the distribution.[39] The second earner profile is calculated as

[37] As demonstrated in Apps and Rees (2018).

[38] The profiles are based on data for all couple income units in the ABS SIH2015-16 with a the primary earner aged from 20 to 64 years, earning a minimum of $12/hour and working for at least 30 hours per week. The sample contains 5481 records .

[39] The Lowess method is applied to obtain a smoothed profile.

the average second earner wage at each primary wage percentile.[40]

Figure 2 about here.

The percentile distribution of the primary earner wage rate is quite flat and virtually linear up to the 80th percentile, but then rises sharply. The degree of inequality between the top percentiles and the rest of the primary earner wage distribution is large, and far outweighs the degree of inequality between primary and average second earner wages. These observations suggest that a far more equitable outcome would have been achieved by raising top tax rates rather than the rates on low and middle wage earners and partnered mothers as second earners.

Against this, it could be argued that raising tax rates on top incomes would cause higher efficiency losses. However, this objection is not supported empirically. In contrast to the wage profiles in Figure 2, hours profiles are relatively flat from around the 20th percentile. As a consequence, primary earner wage elasticities are found to approach zero towards the top percentiles.[41]

We now turn to an empirical analysis of the Australian family tax system to show how using equivalised income as a measure of inequality can support the policy shift just outlined. Section 4.1 focuses on the individual based income tax under the PIT+LITO rate scale, while Section 4.2 identifies the effects of combining the FBT-A system with the PIT+LITO.

4.1 Australian income tax

Personal Income Tax (PIT):

The 2015-16 PIT marginal tax rate (MTR) scale applying to taxable individual income brackets is strictly progressive as shown below.

2015-16 PIT MTR scale:

Bracket	Taxable income bracket	MTR
1	$0 – $18,200	0%
2	$18,201 – $37,000	19%
3	$37,001 – $80,000	32.5%
4	$80,001 – $180,000	37%
5	> $180,000	45%

PIT plus Low Income Tax Offset (PIT+LITO):

The rate scale of the Australian individual-based income tax is set by the PIT scale together with the LITO. The latter provides a tax offset of $445, withdrawn at a rate of 1.5% at an income above $37,000. There are two effects. First, the LITO raises the zero rated threshold from $18,200 to $20,542. Second, the LITO raises the MTR from $37,001 to $66,666 from 32.5% to 34.0%, and thereby introduces an additional taxable income bracket. Note that the rate

[40] For non-participants we use predicted wage rates corrected for selectivity bias.

[41] See Andrienko et al. (2016).

scale is no longer strictly progressive. The LITO reduces the transparency of this outcome.

2015-16 PIT+LITO MTR scale:

Bracket	Taxable income bracket	MTR
1	$\$0 - \$20,542$	0%
2	$\$20,543 - \$37,000$	19%
3	$\$37,001 - \$66,666$	34%
4	$\$66,667 - \$80,000$	32.5%
5	$\$80,001 - \$180,000$	37%
6	$> \$180,000$	45%

The following analysis is based on the PIT+LITO rate scale and draws on the SIH2015-16 data set containing 1055 records for "in-work" families with two children under 15 years.

Table 6 reports the data means of H1 and H2 tax payments on primary and second incomes, followed by the household's total income tax payment, across quintiles of primary income, under this rate scale. H1 and H2 primary incomes are closely matching below the 5th quintile. Consequently, taxes on those incomes are also closely matching below the 5th quintile. In contrast, as shown earlier in Table 3, in each quintile there are significant gaps between second incomes, and so between the taxes on second incomes in Table 6. Overall, H2 households pay significantly higher income taxes in total at any given primary income across most of the distribution of primary income - the overall gap is around 30% to 40%, up to the 5th quintile. The much smaller gap, at around 4%, in quintile 5 can be attributed to the sharp rise in the top primary incomes, as shown in Figure 2, together with the tendency for those in the top percentiles to be in single income households. Second earner tax payments in quintile 5 are almost eight times as large for H2 as for H1 households. It is important to note that, at any given primary income, the gap between H1 and H2 income taxes could be reduced by increasing the progressivity of the marginal rate scale on individual incomes. Thus, a more progressive PIT would reduce both the degree of inequality across the distribution and between genders.

Table 7 presents a matching analysis of tax payments across households ranked by equivalised income. Taxes on E1 primary incomes are substantially higher than those on E2 primary incomes, because of the re-ranking effects of using equivalised joint income. Those on the H1 second incomes are lower than on H2 second incomes, but not sufficiently so as to reverse the overall finding that E1 households are taxed more heavily than E2 households within each quintile, and therefore overall. Thus, changing the basis of ranking households from primary to joint income completely reverses the conclusion on which household types pay more tax.

Figure 3 graphs H1 and H2 total tax payments across primary income, and Figure 4, E1 and E2 total tax payments across equivalised income, to show visually that we obtain opposite results: according to Figure 3, H2 households are more heavily taxed than H1 households, and according to Figure 4, E1 households are more heavily taxed than E2 households.

On distributional grounds, Figure 3 supports individual taxation. Under an individual-based progressive income tax, a two-earner household pays less tax than a single-earner household with the same household income. This at least compensates to some extent for the untaxed status of home production, which makes the single earner household better off than the two-earner household with the same household income. This in turn reflects the fact that joint household income is a poor indicator of household well-being when domestic output is a significant component of household consumption. In contrast, the gap between E1 and E2 taxes shown in Figure 4 would be construed as supporting an argument for preferring joint taxation.

Tables 6 and 7 about here

Figures 3 and 4 about here

Note that nothing changes between Tables 6 and 7 in terms of the actual taxes paid by every household, as determined by the PIT+LITO structure and the actual earnings of the individuals in every household. The apparent change in who pays how much tax is entirely an artifact of the re-ranking brought about by assigning households to quintiles on the basis of equivalised income rather than primary income, and stems from the shift of low to middle income H2 families in quintiles 1 to 3 of primary income into the higher quintiles, and the corresponding shift of higher primary income H1 households into the lower quintiles, as reported in Table 3 and shown graphically in Figure 1. That implies that we should compare the rankings on the basis of their policy relevance. Which better represents the real distribution of household well-being? Our core argument is that once a realistic value is placed on home production, primary income provides the more accurate ranking. Equivalised joint income ignores the existence of household production and therefore gives misleading results. We now turn to an analysis of the effects of shifting towards joint taxation by targeting family payments on joint income.

4.2 PIT+LITO and FTB-A

When FTB-A is combined with the individual based PIT+LITO the overall system becomes one of partial joint taxation. First we describe the structure of the FTB-A system.

FTB-A provides a "Maximum Rate" payment per dependent child per year as follows:

Child under 13 years	$5,412.95
Child aged 13-14 years	$6,825.50

The total of the "Maximum Rate" payment is withdrawn at 20 cents in the dollar on a family income above $51,027 until the remaining payment falls to the “Base Rate + Supplement" per child of $2,230.15. For family incomes above $94,316 this is then withdrawn at 30 cents in the dollar.

With variation according to the age and number of dependent children, the Maximum Rate payments can be classified as "tagged", that is, the payments vary according to needs known to be associated with "observable" characteristics. Families with the same number of children and of the same age, receive the same universal cash payment. A two-child family with both under 13 receives a total payment of $10,825.90. A family with two children, one aged under 13 and other under 14, receives a tagged total payment of $12,238.45.

Combining the FTB-A system with the PIT+LITO introduces dramatic changes in the structure of marginal tax rates. They vary not only with primary and second incomes, but also across joint incomes and, given the variation in the Maximum Rate payments, with the age and number of dependent children. As emphasised previously, "targeting" cash payments on either individual or joint incomes alters effective marginal tax rates and associated intra-marginal lump sums but not the size of the family payments.

To appreciate fully the nature of the change in the structure of marginal tax rates due to withdrawing FTB-A payments on joint income it is useful to take an example. Consider a family with two children, one aged under 13 and the other under 14 years and therefore in receipt of a Maximum Rate payment of $12,238.45. For this family the total "Base Rate + Supplement" is $4,460.3. The Maximum Rate net of the "Base Rate + Supplement" is therefore $9190.7.

Now consider a single income family with a primary income of $75,000. This income places the family well below the upper threshold limit of quintile 1 of a ranking defined on household income, and therefore of the equivalised income ranking in Table 4. The primary earner pays income taxes at PIT+LITO marginal rates up to $51,027. From $51,028 to $75,000, his marginal tax rate is not the PIT+LITO rate of 34%, but 54%.

Now suppose the second partner goes out to work for $70,000. The first point to note is that the family is shifted from quintile 1 towards quintile 4 of an equivalised income ranking because their joint incomes equate to an equivalised income that is just below that of the lower threshold of quintile 4. The family is now represented as being well within the top half of the distribution of income. Note that this large shift is due also to the very flat profile of primary incomes up to 80th percentile, as discussed above.

The second earner faces a marginal tax rate of 20% on the first dollar earned. She is denied a zero-rated threshold because when she starts working the family's joint income exceeds the lower limit for the 20% withdrawal rate of the Maximum Rate payment. Her marginal tax rate is therefore 20%. She continues to face an MTR that is 20% above the PIT+LITO rate until the family income reaches the income level at which the Maximum Rate is fully withdrawn. At an income of $24,316 ($94,316 - $75,000) her MTR rises by 30% and does not return to the PIT+LITO scale until the Base Rate + Supplement of $4460.3 is fully withdrawn.[42]

Table 8 shows the impact of the FTB-A system on H1 and H2 total income taxes across primary income. Net FTB-A in the first row of each panel of the table is calculated as the "Maximum Rate" net of the implicit tax imposed by shifting towards a less progressive rate scale and towards joint taxation. The H1 and H2 "total tax" figures are calculated by subtracting net FTB-A payments from the household income taxes reported in Table 6. The quintile profiles are

[42] For further numerical examples, see Apps (2017).

shown graphically in Figure 5. We see that the FTB-A system widens the tax gap between H1 and H2 households across low to middle income levels. This is evident when we compare Figures 3 and 5. Thus instead of moving toward a more progressive rate scale on individual incomes which, as noted previously, would reduce the gap between taxes paid by families with the same primary income, the new MTR scale resulting from the FTB-A system increases that gap.

This result reflects the way in which introducing joint income-targeted benefits widens the net-of-tax gender pay gap, by raising effective marginal tax rates on second earners. When the Government is simultaneously reducing the progressivity of the rate scale, either by lowering top rates or by raising the bracket limits to which they apply, there is also a shift in the overall tax burden towards the lower and middle ranges of primary income. The Government has the option of funding family payments from higher PIT rates on households with and without children across the entire income distribution and, given the wage profile in Figure 7, one might expect to see a contribution from higher tax rates on top incomes.

Table 9 reports the impact of FTB-A on E1 and E2 total taxes across equivalised income. Net FTB-A is again calculated as the maximum rate net of the implicit additional tax due to shifting towards joint taxation and a less progressive rate scale. Total taxes are calculated by subtracting net FTB-A payments from household income taxes reported in Table 7.

Tables 8 and 9 about here

Figures 5 and 6 about here

When compared with Table 8, we see that an equivalised income ranking gives very different results. The gains from FTB-A are concentrated almost entirely in quintiles 1 and 2 and reduce the larger absolute gaps between E1 and E2 tax burdens reported for those quintiles in Table 7. The new distribution of E1 and E2 tax burdens is shown graphically in Figure 6. When compared with Figure 4, we see not only significantly lower but also more equal E1 and E2 total taxes in those quintiles, an outcome which will be viewed mistakenly as a gain in equity.

5 Conclusions

In this paper we have presented a critique of the standard methods of using equivalence scales to make across-household welfare comparisons. The basis of this critique is the argument that they produce misleading results in support of tax/transfer policies which disadvantage low to middle income households and the women in those households as second earners. In particular they supply arguments in favour of joint taxation and the withdrawal of child benefits on the basis of joint income, which make working mothers significantly worse off and so create disincentives to female labour force participation.

More constructively, we have proposed alternatives to inequality measures based on household income deflated by an equivalence scale. These have more relevance to the analysis of public policies, particularly those relating to taxes and transfers. From a theoretical point of view, we should develop measures based on indirect utility or expenditure functions defined on truly exogenous variables, and which reflect significant differences in the preference structures and constraints that different households face. In particular we should distinguish between groups of households in different stages of the family life cycle, and also those with particular characteristics such as the presence of disabilities or being a lone parent, that are amenable to being tagged. To impose homogeneity in the measurement of well-being across all such groups by an equivalisation procedure obscures important issues, such as the trade-offs that tax/transfer policy has to make between these groups of households. It is essential to disaggregate into groupings across which preference structures and constraints differ significantly.

We fully recognise that such disaggregation requires data sets that are currently not available, though this is at least in part due to the crowding-out effect of the routine procedures that standard measures based upon equivalised household incomes offer. As an interim solution we propose, for family households in which household production is a significant form of time use, primary earner income as a better, readily available measure of household well-being. A large part of the paper has been devoted to making the argument, in the concrete policy setting of the Australian family tax system, that this is a less distortionary basis for household welfare rankings and public policy more generally.

It appears to be the case that the gender wage gap in terms of pre-tax/transfer wage rates has improved somewhat over the past few decades, but unfortunately inappropriate policies supported by an overly simplistic economic methodology have had at least to some extent a countervailing effect on net wage rates, which are after all more fundamental determinants of household well-being than household income, an endogenous variable.[43] To the extent that the position of women within the household is influenced by their outside options, which seems to be a characteristic of almost all modern work on models of the family household, it cannot have been improved by policies that weaken them.

Appendix 1

Equivalised income measures: the theory

Here we set out the theoretical derivation of the equivalisation procedure, following Deaton and Muellbauer.[44] Assume that every household $h = 1, ..., H$ possesses a utility function of the form

$$u_h = u(x^h, a^h) \tag{10}$$

where $u(.)$ is the same function for all households and has the properties of the individual utility function in the standard analysis of consumer demand, and

[43] See for example the empirical study by Bick et al. (2018), which identifies taxation as the driving force behind developments in female labour supply over the past few decades.

[44] The notation is ours.

x^h is a vector of consumption goods bought on markets at a price vector p, assumed to be the same for all households.[45] The key assumption here is that the vector of household characteristics a^h captures everything that would make two households with the same income[46] y_h (and facing the same price vector) choose different consumption vectors and, moreover, that these are observable to the analyst and so can be controlled for. We will refer to a^h as the household's *type*. These characteristics may take any form but the elements of a^h most used in practice are demographic variables such as the number of household members and their ages.

In the usual way, we can derive the household's indirect utility function $v(p, y_h, a^h) = u(\tilde{x}^h, a^h)$, where $\tilde{x}^h$ is the optimal consumption vector at p, y_h, a^h, and its expenditure function $e(p, a^h, u^h)$, giving the minimum expenditure required to achieve the utility u_h. Again the functions $v(.)$ and $e(.)$ are the same for all households. This implies that households with the same income and type have identical utilities and, holding income constant, the only thing that causes variation in demands and utilities is variation in type. This naturally suggests construction of an index number for household utility along the following lines.

We can choose arbitrarily, but without loss of generality, a reference household type, denoted by a^0, for example a household with a single individual. Recall that given the usual assumptions of consumer theory we can always define a *money metric at constant prices* for utility, that is, we can take as our utility measure the minimised expenditure required by type a^h at the given price vector p to achieve a given utility level. So for any p, a^h we can write $u_h \equiv e(p, a^h, v(p, y_h, a^h))$.

We now define the equivalence scale index number as

$$\mu_h = \frac{e(p, a^h, v(p, y_h, a^h))}{e(p, a^0, v(p, y_h, a^h))} \tag{11}$$

The denominator is the amount of expenditure required by the type a^0 to achieve the utility level $v(p, y_h, a^h)$, the numerator is the amount of expenditure required by type a^h to reach the same utility level. Thus $\mu_h \gtreqless 1$ according as $e(p, a^h, v(p, y_h, a^h)) \gtreqless e(p, a^0, v(p, y_h, a^h))$, so that for example a household that requires twice the expenditure of the reference household to achieve the same given level of utility has an index value of 2.

Note that:

$$y_h \equiv e(p, a^h, v(p, y_h, a^h)) \tag{12}$$

by definition of the expenditure function, while setting $a^h = a^0$ and the required utility level at $v(p, y_h, a^h)$ gives

$$u_h = e(p, a^0, v(p, y_h, a^h)) \tag{13}$$

[45] Deaton and Muellbauer (1980) show that price vectors can differ across households, as long as they are fully observable by the analyst. In that case a price index number can be calculated that becomes part of the equivalisation index.

[46] Note that since the model has no saving, income and expenditure are equal. Deaton and Muellbauer show how this can be generalised.

at the reference type. Thus, using (2) we can write utility as

$$u_h = \frac{y_h}{\mu_h} \tag{14}$$

which provides the theoretical basis for the equivalisation procedure.

Appendix 2

The Household Exchange Model

Households consist of a primary and a second earner. Both divide their time between market labour supply L, leisure l, and time a spent in production of a household good z, that they both consume.[47] They have market wage rates of w_{ih}, $i = 1, 2$, with $w_{1h} \geq w_{2h}$, and $h = 1, 2, ..., H$ denotes the household.

On the production side, the household chooses its time allocation efficiently by solving the problem[48]

$$\min_{a_{ih}, b_h} \sum_{i=1}^{2} w_{ih} a_{ih} + p_h b_h \quad \text{s.t. } z_h \geq f(k_{1h} a_{1h}, k_{2h} a_{2h}, q_h b_h) \quad h = 1, 2, ..., H \tag{15}$$

The production function $f(.)$, is identical across households, linear homogeneous in effective labour supplies and strictly quasiconcave.

By setting $z_h = 1$ the solution to this problem yields unit demand functions

$$a_{ih}^{(1)} = a(w_{1h}, w_{2h}, p_h, k_{1h}, k_{2h}, q_h); \; b_h^{(1)} = b(w_{1h}, w_{2h}, p_h, k_{1h}, k_{2h}, q_h) \tag{16}$$

and a unit cost function $c(w_{1h}, w_{2h}, p_h, k_{1h}, k_{2h}, q_h)$ independent of the level of output. This defines an implicit price of the domestic good, denoted by

$$\pi_h = c(w_{1h}, w_{2h}, p_h, k_{1h}, k_{2h}, q_h) = \sum_i w_{ih} a_{ih}^{(1)} + p_h b_h^{(1)} \tag{17}$$

The Envelope Theorem gives, for $i = 1, 2, \; h = 1, ..., H$:

$$\frac{\partial \pi_h}{\partial w_{ih}} = a_{ih}^{(1)}; \frac{\partial \pi_h}{\partial p_h} = b_h^{(1)}; \frac{\partial \pi_h}{\partial k_{ih}} = -\alpha_h f_i a_{ih}^{(1)}; \frac{\partial \pi_h}{\partial q_h} = -\alpha_h f_3 b_h^{(1)} \tag{18}$$

where $\alpha_h > 0$ are Lagrange multipliers. Because of the linear homogeneity assumption we can write the input demand functions as

$$a_{ij} = a_{ih}^{(1)} z_h; \; b_h = b_h^{(1)} z_h \tag{19}$$

while $a_{ih}^{(1)} z_{jh}$ denotes the amount of time i spends in producing the amount of z consumed by individual j, $i, j = 1, 2$ and $\sum_i z_{ih} = z_h$.

[47] We treat z as a private good rather than as a household publlc good. It is straightforward to extend the model to incorporate this, using Paul Samuelson's analysis of optimal public good choices, but that does not add anything of interest given the focus of this paper.

[48] Details of this solution are given in the appendix.

The individual time constraints are:

$$a_{ih} + l_{ih} + L_{ih} = T \tag{20}$$

The constant returns to scale assumption also implies that there is a separation between production and consumption, which greatly simplifies the analysis.

Turning to the consumption side, the individual utility functions are:

$$u_{ih} = u(x_{ih}, z_{ih}, l_{ih}) \quad i = 1, 2 \tag{21}$$

where x_{ih} denotes consumption of a composite market good and preferences are identical across all individuals and households, which allows interpersonal comparisons of utility.

We assume that the household is fully rational in that it values each earner's time consistently at their outside market wage. In particular it prices individual leisures at their corresponding market wage and applies the implicit price π_h as the opportunity cost of each individual's consumption of the household good. This implies that we can represent the individual choice problems in terms of their *full income budget constraints*:

$$\max_{x_{ih}, z_{ih}, L_{ih}} u(x_{ih}, z_{ih}, L_{ih}) \text{ s.t. } x_{ih} + \pi_h z_{ih} + w_{ih} l_{ih} \leq w_{ih} T \quad i = 1, 2, \quad h = 1, 2, ..., H \tag{22}$$

Equivalently, we could write the individual budget constraints in terms of income and expenditure.

$$x_{ih} + (w_{jh} h_{jh}^{(1)} + p_h b_h^{(1)}) z_{ih} \leq w_{ih}(l_{ih} + h_{ih}^{(1)} z_{jh}) \quad i, j = 1, 2, \ i \neq j \tag{23}$$

We derive these expressions by substituting for $T = l_{ih} + L_{ih} + a_{ih}$ and using $\pi_h = \sum_i w_{ih} h_{ih}^{(1)} + p_h b_h^{(1)}$. The left hand sides of these equations give the expenditure on i's consumption of the market good and the cost of the inputs required for producing i's consumption of the household good, other than its own input, and the right hand sides give the sums of i's market wage earnings and the implicit payment from the partner j for the time i spends in producing j's consumption of the household good, $i, j = 1, 2$, $i \neq n$. This recognises that each individual has not one but two sources of income: as well as labour market earnings there is the implicit income from participating in producing the partner's share of the household good.

Summing these two constraints gives:

$$\sum_i x_{ih} + p_h b_h = \sum_i w_{ih} l_{ih} \tag{24}$$

which is the household's "balance of payments" constraint: the cost of its "imports" in the form of the market consumption good and market input into household production must be covered by the value of its "exports", its market labour supplies. The within-household transactions with respect to the non-traded good z_h of course cancel out in the aggregate. This emphasises the view of the household as a small economy and highlights the limitations of models that ignore its non-traded good and essentially set the value of its GDP equal to its exports.

That the second earner may have a cost of consumption of market goods greater then her market income should not be interpreted as implying that she receives a *lump sum* transfer from the primary earner. To prove the existence of a lump sum transfer in this case, it is necessary to show that her consumption is greater than the amount of her market income plus the value of her contribution to 1's consumption of the market good, net of the cost of the bought in market input required to produce her consumption of the domestic good i.e.:

$$x_{2h} > w_{2h}(l_{2h} + h^{(1)}_{2h} z_{1h}) - p_h b^{(1)}_h z_{2h} > 0 \quad (25)$$

For example, a high wage second earner who supplies little labour time to the market and provides a relatively large amount of the household good to her partner, with relatively little consumption of it herself, could have in a no-transfer equilibrium a large excess of market consumption over her own market labour income. Essentially, in this case the second earner is trading her time in producing the household good for her partner's time in earning the market good. This again demonstrates the limitations of household models that ignore household production.

Solving the individual choice problems yields individual demand functions $x_{ih}(w_{ih}\pi_h)$, $z_{ih}(w_{ih}, \pi_h)$, $l_{ih}(w_{ih}, \pi_h)$ with all the standard properties, as well as individual indirect utility functions $v_{ih}(w_{ih}, \pi_h)$ and expenditure functions $e_{ih}(w_{ih}, \pi_h, u_{ih})$. These functions and their properties form the basis for the discussion of inequality in the text of the paper. Since our focus there will be on the indirect utility function, we present here its derivatives with respect to its exogenous determinants.

Because of the two-stage form of our analysis the derivatives of $v_{ih}(w_{ih}, \pi_h)$ derived from the household consumption decisions are deceptively simple:

$$\frac{\partial v_{ih}}{\partial w_{ih}} = \lambda_{ih}(T - l_{ih}); \quad \frac{\partial v_{ih}}{\partial \pi_h} = -\lambda_{ih} z_{ih} \quad i = 1, 2 \quad (26)$$

where λ_{ih} is the marginal utility of i's full income. However, this suppresses the fact that π_h is a function of all the variables in the vector $[w_{1h}, w_{2h}, p_h, k_{1h}, k_{2h}, q_h]$. Thus the full derivatives of the indirect utility function with respect to the truly exogenous determinants of utility are both more complex and more interesting:

$$\frac{\partial v_{ih}}{\partial w_{ih}} = \lambda_{ih}[(T - (l_{ih} + z_{ih} h^{(1)}_{ih})] = \lambda_{ih}[l_{ih} + z_{jh} h^{(1)}_{ih}] > 0 \quad i, j = 1, 2, \ i \neq j \quad (27)$$

This derivative shows that an increase in the wage changes utility proportionally to not just market labour supply but also to the time given both to this and to the production of the household good for consumption of the other individual in the household, because of the implicit trade relationship. For example an increase in her wage makes the second earner better off even if she works a negligible amount in the market because it raises the implicit return to her household labour supply, given that the household rationally values her time at her market wage rate.

$$\frac{\partial v_{ih}}{\partial w_{jh}} = -\lambda_{ih} z_{ih} h_{jh}^{(1)} < 0 \quad i,j = 1,2, \; i \neq j \tag{28}$$

An increase in one individual's wage however has a negative effect on the utility of the other, because it raises the price of the household good. This utility effect is therefore proportional to the amount of the household good consumed by the individual whose wage has not risen, (a standard Roy's Identity effect), as well as to the time input per unit of output of the individual whose wage has risen (the household production effect).

$$\frac{\partial v_{ih}}{\partial p_h} = -\lambda_{ih} z_{ih} b_h^{(1)} < 0 \quad i = 1,2 \tag{29}$$

Of course an increase in price of the bought in input increases the implicit price of the household good and so makes both individuals worse off, to an extent dependent on their consumption of the household good (Roy's Identity again) and the amount of the good used per unit of output of the household good (household production effect).

$$\frac{\partial v_{ih}}{\partial k_{ih}} = \lambda_{ih} \alpha_h f_i z_{ih} a_{ih}^{(1)} > 0 \quad i = 1,2 \tag{30}$$

$$\frac{\partial v_{ih}}{\partial k_{jh}} = \lambda_{ih} z_{ih} \alpha_h f_j a_{jh}^{(1)} > 0 \quad i,j = 1,2, \; i \neq j \tag{31}$$

$$\frac{\partial v_{ih}}{\partial q_h} = \lambda_{ih} \alpha_h f_3 z_{ih} b_h^{(1)} > 0 \quad i = 1,2 \tag{32}$$

Increases in each of the productivity variables lowers the price of the household good and so must make each individual in the household better off. The size of the effect varies positively with the marginal product of the input concerned and the amount of it used per unit of the household good, as well as with the amount of the good the individual consumes (yet again Roy's Identity).

These derivatives are all perfectly intuitive. Their main aim is not to consider comparative statics effects on any one household in isolation, but rather to consider *what happens to the standard of living of households as we move through the joint population distribution of this vector of variables and observe how this relates to the measurement of household inequality.*

References

[1] Alesina, A., A. Ichino and L. Karabarbounis (2011). Gender based taxation and the division of family chores. *American Economic Journal: Economic Policy,* 3 (2), 1-40.

[2] Andrienko, Y., P. Apps and R. Rees (2016). Optimal taxation and top incomes. *International Tax and Public Finance* 23 (6), 981-1003.

[3] Apps, P. (1982). Institutional inequality and tax incidence. *Journal of Public Economics.* 18(2), 217-42.

[4] Apps, P. (2017). Gender equity in the tax-transfer system for fiscal sustainability. Ch 3 in M Stewart (ed) *Tax, Social Policy and Gender.* 69-98, ANU Press, Canberra.

[5] Apps, P., N.V. Long and R. Rees (2014). Optimal piecewise linear income taxation. *Journal of Public Economic Theory*, 16 (4), 523-545.

[6] Apps, P. and R. Rees (1988). Taxation and the household. *Journal of Public Economics* 35 (3), 355–69.

[7] Apps, P. and R. Rees (2009). *Public Economics and the Household.* Cambridge: Cambridge University Press.

[8] Apps, P. and R. Rees (2011). Household time use, inequality and taxation. *Household Economic Behaviors.* Ch 3 in J Molina (ed), Springer.

[9] Apps, P. and R. Rees (2018). Optimal family taxation and income inequality. *International Tax and Public Finance.* 25 (5), 1093-1128.

[10] Atkinson, A. (1970). On the measurement of inequality. *Journal of Economic Theory, 2, 244-63.*

[11] Atkinson, A„ C Leventi, B. Nolan, H. Sutherland and I Tasseva (2017), "Reducing poverty and inequality through tax benefit reform and the minimum wage: The UK as a case study. *Journal of Economic Inequality*, 15, 20, 303-323.

[12] Basu, K. (2006). Gender and say: A model of household behaviour with endogenously determined balance of power. *The Economic Journal.* 116, 558-580.

[13] Bergmann, B. (1971). The effects on white incomes of discrimination in employment. *Journal of Political Economy.* 79, 294-313.

[14] Bergstrom, T, 1993, A survey of theories of the family, in *Handbook of population and family economics, M R Rosenzweig and O Stark, vol 1, ch 2, 21-79.*

[15] Bick, A., B. Brueggemann, N. Fuchs-Schuendeln and H. Paule-Paludkiewcz (2018). Long-term changes in married couples' labor supply and taxes: Evidence from the US and Europe since the 1980's. *CESifo WP* 7267.

[16] Chiappori. P.-A. (1988). Rational household labor supply. *Econometrica,* 56, 63–90.

[17] Deaton, A. and J. Muellbauer (1980). *Economics of consumer behavior.* Cambridge University Press, Cambridge.

[18] Gayle G-L and A Shephard, 2019, Optimal taxation, marriage, home production and family labor supply", *Econometrica*, 87, 1, 291-326.

[19] Konrad K. and K. Lommerud (2000). The bargaining family revisited. *Canadian Journal of Economics.* 33 (2), 471-87.

[20] LaLumia, S. (2008). The effects of joint taxation of married couples on labor Supply and non-wage Income. *Journal of Public Economics.* 92 (7), 1698-1791.

[21] Lundberg, S., R. Pollak and J. Stearns (2016). Family Inequality: Diverging patterns in marriage, cohabitation and childbearing. *Journal of Economic Perspectives.* 30 (2), 79-102.

[22] Lundberg, S., R. Pollak and T. Wales (1997). Do husbands and wives pool their resources? Evidence from the UK child beneft. *Journal of Human Resourcs.* 232(3), 463-80.

[23] Mirrlees J. (1971). An exploration in the theory of optimum Income taxation, *Review of Economic Studies*, 38, 175-208.

[24] OECD (2019), Under pressure: The squeezed middle class, OECD Publishing, Paris.

[25] Productivity Commission (2018), Rising inequality? A stocktake of the evidence, Productivity Commission Report, Productivity Commision, Australia.

[26] Pollak, R. and T. Wales (1979). Welfare comparisons and equivalence scales. *American Economic Review*, 69, 316-21.

[27] Pollak, R. and M. Wachter (1975). The relevance of the household production function and its implications for the allocation of time. *Journal of Political Economy.* 68, 349-59.

[28] Sila, U. and V. Dugain (2019). Income, wealth and earnings inequality in Australia: Evidence from the Hilda survey. WP No. 1538. Economics Department, OECD, Paris.

Table 1 Quintile distribution of primary income and participation rates

Quintile	1	2	3	4	5	All
Prim. income $pa	50025	74620	95825	123166	248023	118332
2nd income $pa	21641	28222	37255	40618	57277	37003
Second earner participation rates						
Non-participants %	0.29	0.27	0.24	0.28	0.25	0.27
Part time %	0.43	0.42	0.44	0.48	0.49	0.45
Full time %	0.28	0.31	0.32	0.24	0.26	0.28

Table 2 Quintile distribution of equivalised income and participation rates

Quintile	1	2	3	4	5	All
Prim. income $pa	56470	75722	98689	121625	239156	118332
2nd income $pa	8630	22007	30761	47071	76546	37003
Second earner participation rates						
Non-participant %	0.54	0.30	0.26	0.11	0.12	0.27
Part time %	0.34	0.48	0.46	0.51	0.47	0.45
Full time %	0.12	0.22	0.28	0.38	0.41	0.28

Table 3 Primary and 2nd incomes by primary income, $pa

Quintile		1	2	3	4	5	All
H1	Prim. income	49782	74381	95897	123786	266330	122035
	2nd income	10887	11381	13699	16395	22106	14894
	H'hold income	**60669**	**85762**	**109596**	**140181**	**288436**	**136929**
H2	Prim. income	50268	74859	95753	122546	229716	114628
	2nd income	32395	45063	60811	64841	92448	59111
	H'hold income	**82663**	**119922**	**156564**	**187387**	**322164**	**173739**

Table 4 Primary and 2nd incomes by equivalised income, $pa

Quintile		1	2	3	4	5	All
E1	Prim. Income	60725	88105	112942	139134	284008	136983
	2nd income	412	9165	15303	28995	50083	20792
	H'hold income	**61137**	**97270**	**128245**	**168129**	**334091**	**157660**
E2	Prim. income	52215	63338	84436	104115	194305	99682
	2nd income	16848	34848	46219	65148	103009	53214
	H'hold income	**69063**	**98186**	**130655**	**169263**	**297313**	**152623**

Table 5 Distribution of H1 and H2 households by equivalised income

Quintile	1	2	3	4	5	All
H1 quintile split %	76	57	46	36	28	50
H2 quintile split %	24	43	54	64	72	50

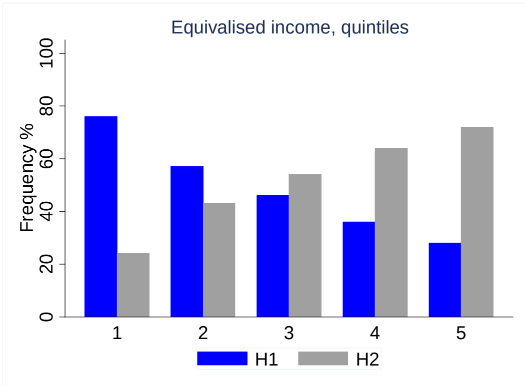

Figure 1 Quintile distribution of H1 and H2 households

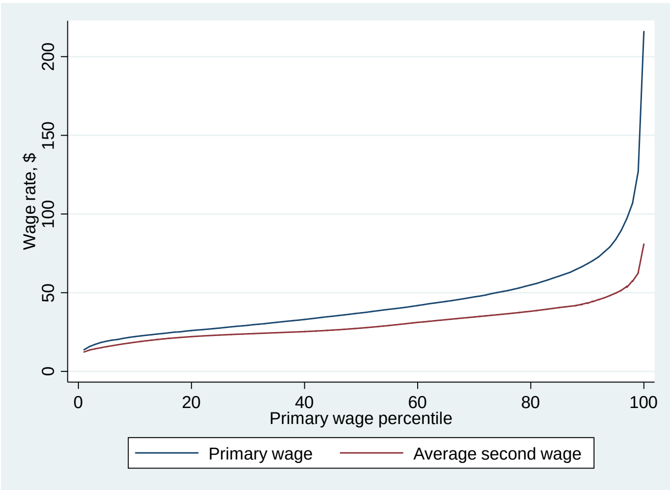

Figure 2 Percentile primary and average second wage distributions

Table 6 Income tax payments, $pa: H1 and H2 by primary income

Quintile		1	2	3	4	5	All
H1	Prim. Income tax	7487	15531	22707	32767	92433	34185
	2nd income tax	706	640	979	2076	2993	1479
	H'hold income tax	**8193**	**16171**	**23688**	**34843**	**95426**	**35664**
H2	Prim. income tax	7526	15600	22873	32507	76244	30950
	2nd income tax	3100	6731	10195	12860	23108	11199
	H'hold income tax	**10626**	**22331**	**33068**	**45367**	**99352**	**42149**

Table 7 Income tax payments, $pa: E1 and E2 by equivalised income

Quintile		1	2	3	4	5	All
E1	Prim. Income tax	11087	20400	29166	38127	99932	39742
	2nd income tax	0	333	1158	3411	10388	3058
	H'hold income tax	**11087**	**20733**	**30324**	**41538**	**110329**	**42800**
E2	Prim. income tax	8054	11795	18852	25240	61047	24998
	2nd income tax	848	3109	6409	12181	26662	9842
	H'hold income tax	**8902**	**14904**	**25260**	**37421**	**87709**	**34839**

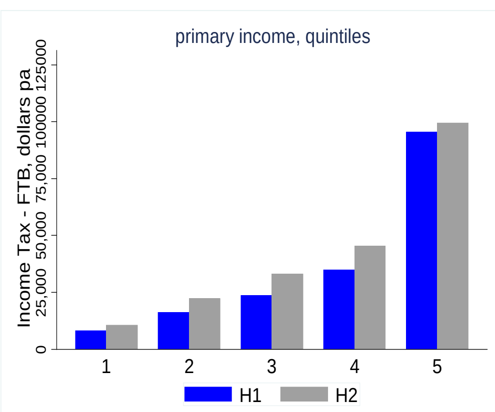

Figure 3 H1&H2 income tax payments

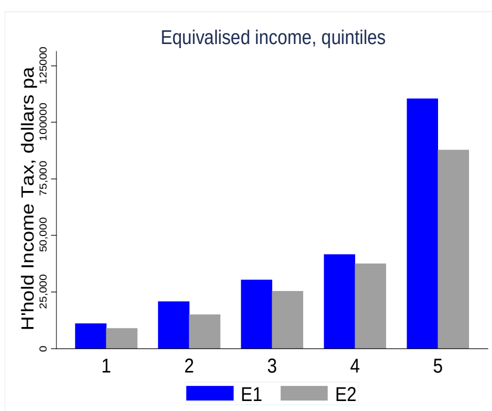

Figure 4 E1&E2 income tax payments

Table 8 Income tax – FTB-A payments, $pa: H1 and H2 by primary income

Quintile		1	2	3	4	5	All
H1	FTB-A	8577	4791	2049	43	0	3093
	Total tax	**-384**	**11380**	**21639**	**34800**	**95426**	**32571**
H2	FTB-A	5406	895	99	0	0	1280
	Total tax	**5220**	**21436**	**32969**	**45367**	**99352**	**40869**

Table 9 Income tax – FTB-A payments, $pa: E1 and E2 by equivalised income

Quintile		1	2	3	4	5	All
E1	FTB-A	8428	3027	46	0	0	2300
	Total tax	**2659**	**17706**	**30278**	**41538**	**110329**	**40500**
E2	FTB-A	7087	3167	67	30	0	2070
	Total tax	**1825**	**11737**	**25193**	**37391**	**87709**	**32771**

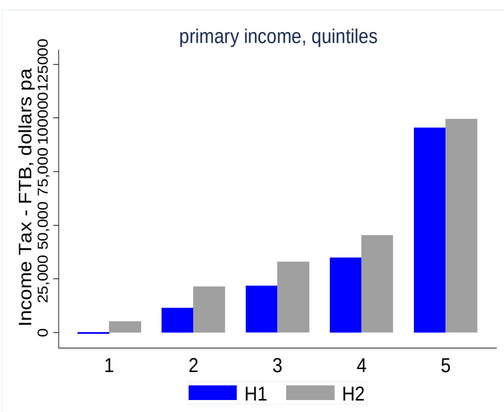

Figure 5 H1&H2 Income tax – FTB-A

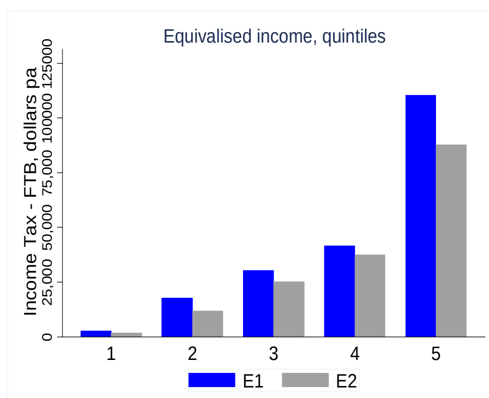

Figure 6 E1&E2 Income tax – FTB-A