Individual Welfare Analysis: What's the Role of Intra-Family Preference Heterogeneity?*

Tim Obermeier

Institute for Fiscal Studies

February 15, 2022

Abstract

How well does household income reflect the welfare of the individuals living within the household? To answer this question, I build a structural model of family decision-making and the marriage market. Household income does not take unequal sharing of consumption within families, time use (leisure and housework) and individuals' preferences into account. In particular, partners value the household-level public good differently if their preferences differ. I estimate the model based on time use data from the UK, using the marriage market equilibrium to pin down the unobserved joint distribution of preferences of partners. I use the estimated model to contrast poverty in household income and welfare. The main finding is that only 61% of welfare-poor individuals are also poor in terms of household income, highlighting that a substantial fraction of welfare-poor individuals is missed by the conventional focus on income. Preference heterogeneity plays an important role for this result. Taking only unequal sharing and time use - but not preferences - into account, at most 81% of welfare-poor individuals can be detected, which leaves 19% unidentified. Finally, to illustrate the policy relevance of individual welfare measures, I study how minimum wage increases affect welfare-poverty in this framework.

Keywords: Individual Welfare, Preference Heterogeneity, Inequality, Marriage Market, Intra-Household Inequality, Minimum Wage

JEL classification: E21, I32, D13, D63

*I thank Richard Blundell, Christian Bredemeier, Monica Costa Dias, Bram De Rock, Falko Jüßen, Jan Gravert, Anne Hannusch, Karen Kopecky, Valérie Lechene, Krishna Pendakur, Franck Portier, Michèle Tertilt and John Van Reenen, as well as seminar participants at IFS, UCL, the University of Wuppertal and various other universities for many very helpful comments. Support by the state of Baden-Württemberg through bwHPC is gratefully acknowledged. Email: `tim.obermeier@ifs.org.uk`.

1 Introduction

Studying inequality has been a central task of economists in recent years as concerns about inequality have risen in many countries. Many studies focus on inequality in household income, sometimes converted to consumption using adult equivalence scales (see e.g. Atkinson and Bourguignon (2014), Heathcote, Storesletten, and Violante (2010)). While the focus on household income - and the implicit assumption that all household members are equally well off - is often necessary for practical reasons, it also imposes several restrictions which could lead to an incomplete picture of inequality. There are at least three reasons why *household* income does not fully capture the well-being of the *individuals* living within the family. First, consumption can be shared unequally within households, so that household income may not reflect the consumption level of the individual. This has been a focus of the large literature on collective household models, which typically finds that equal sharing of consumption within families is rejected by the data (e.g. Lise and Seitz (2011), Dunbar, Lewbel, and Pendakur (2013), Cherchye et al. (2015), Lise and Yamada (2019)).[1] Second, household income ignores the allocation of time (in terms of leisure, housework and market hours), which is an important determinant of welfare. For example, individuals who have little leisure are less well-off in terms of welfare than it might appear based on just looking at their household income.[2] Third, household income does not take individuals' preferences into account, i.e. how much they value different consumption goods and time use, which is the final element required for drawing conclusions about welfare. These considerations are important both for our understanding of inequality and for the targeting of redistributive policies.

As a simple example, consider a household with two members, A (Alex) and B (Bertie), with total monthly income of £2000. Alex and Bertie consume two goods, a private and a public good, and choose between working and leisure. Suppose we observe that Alex spends £400 on private consumption and Bertie spends £600. In addition, Alex works more and only has 20 hours of leisure per week while Bertie has 25 hours of leisure. Finally, the couple spends an important fraction of their budget, £1000, on the public good, which is consumed equally by both partners.

The example illustrates that the welfare of Alex and Bertie depends on more than their household income. It also highlights that knowing more about preferences is essential to determine who is better off. If we are willing to assume that both partners have equal preferences, Alex is worse off than Bertie, since they have a lower level of private consumption and leisure and both partners benefit equally from the public good. However, any welfare ranking is possible once we allow for preference heterogeneity. The fact that Alex has lower consumption and leisure is not necessarily a sign that they are in a disadvantaged position within the household. Instead, an equally plausible scenario is that their priorities are different - they might care less about these

[1] For example, Lise and Seitz (2011) estimate a collective model and find that intra-household inequality accounts for 25% of total consumption inequality in their sample of childless British households.

[2] As a result, some studies have taken leisure into account when computing welfare measures (e.g. Cherchye et al. (2018), Lise and Seitz (2011)).

goods and value the public consumption good more strongly. Viewed from this angle, Alex could be better off than Bertie or there might be no inequality between the two at all. If the goal of the analysis is to determine which individuals in the population should be considered as poor, either Alex or Bertie - or both - could be poor in terms of their welfare (i.e. be among the least well-off in the population), even when the household is above the poverty line for income.

To study these issues, I build a structural model of household decision-making, preference heterogeneity and the marriage market and estimate the model using data from the UK. I then use the estimated model for an individual-level welfare analysis, which allows to assess to what extent welfare-poor individuals are also poor in terms of household income. In addition, I study the policy implications of taking an individual-based view of welfare by simulating a minimum wage increase and analysing its impact on welfare-poverty.

The model builds on the example of Alex and Bertie and considers an economy with both consumption and time use which is comparable to standard setups used in the collective model literature (see e.g. Lise and Yamada (2019)). Households choose between private and public expenditure and the allocation of time between market work, leisure and domestic time. The public good within the household is 'produced' using public expenditure and domestic time of each household member. The novelty in my model is adding multi-dimensional preference heterogeneity and a search-based marriage market.[3] The preferences of each individual are drawn from a population distribution, so that there is heterogeneity within the groups of men and women. The marriage market endogenises the formation of couples.

This framework allows to address the empirical challenge that preferences are unobserved in the data, as there are no direct measures of the type of preferences needed in collective household models (i.e. the valuation of consumption, leisure and the public good). Unobserved preference heterogeneity is difficult to incorporate into state-of-the-art microeconometric estimation strategies, such as Lise and Yamada (2019), as preference heterogeneity needs to be separated from unobserved differences in the extent of sharing within the household.[4] To address this issue, I exploit the marriage market equilibrium. The key idea behind the model is that data on the variability of choices among singles, together with a theory of the marriage market, determines the joint distribution of unobserved preferences of partners. This approach only requires to estimate the parameters of the *population* distribution of preferences. Given the population distribution, the marriage market determines who gets married to whom and who stays single. This avoids having to estimate how *correlated* preferences are between partners, as these correlations are endogenous. The parameters of the population distribution can be estimated to match the variability of time use choices among singles. This approach is related to what Chiappori and Meghir (2015) call

[3] Throughout the paper, marriage will refer to both married and cohabiting couples in the data.

[4] As an example, suppose we observe that a woman consumes more than her partner. This could either by rationalised as unobserved preference heterogeneity (i.e. her having a higher consumption preference) or as an unobserved difference in 'bargaining power' (preferences being equal but the couple sharing consumption unequally). As a result, allowing for unobserved preference heterogeneity in standard estimation strategies is difficult, as it requires disentangling the two.

'identifying from market equilibrium', which has been applied in other contexts.[5]

To estimate the model, I use data from the UK Time Use Survey (UKTUS), using the waves 2000/1 and 2014/15 and the Longitudinal Household Survey (Understanding Society/USOC). The UKTUS contains detailed time diaries from which allow to compute time spent on market work, leisure and so-called home production (chores and caring for children). Variability in work hours, leisure and home production choices among singles serve as an indicator of preference heterogeneity. Work hours choices are a proxy for the consumption preference. The parameters of the preference distribution are estimated so that the distribution of singles from the model replicates the variability of choices that is observed in the data. In addition, I use data from USOC on wages, as the UKTUS data contains limited information on earnings and wages.

I use the estimated model for an individual-level poverty analysis. To measure individual welfare, the model allows to compute the *Money-Metric Welfare Index* (MMWI, Chiappori and Meghir (2015)). This converts an individuals' utility from consumption, leisure and the home good into a monetary index that can be compared across individuals with heterogenous preferences. Poverty, both for household income and welfare, is defined as being in the lowest 20% of the corresponding distribution. This holds the number of poor individuals constant across welfare measures, but not all of those are income-poor are also welfare-poor (and vice versa). In particular, welfare-poor individuals may live in households in any position of the income distribution. For example, in the 4th decile of household income, where households are above the poverty line for income, around 15% of individuals are poor on an individual level.

The main finding is that defining poverty based on household income only identifies 61% of individuals who are welfare-poor. In other words, 39% of welfare-poor individuals would be missed by focusing on household income. This demonstrates that taking a broader view of welfare and including (1) unequal sharing within families (2) time use and (3) preferences can lead to different conclusions than a conventional analysis based on household income. There are two groups of people who are welfare-poor but not income-poor. The first are those individuals who have low "bargaining power" within the family and only reach low levels of welfare even though the household is not income-poor. The second are those who live in families with low economies of scale. Economies of scale (the extent to which partners can jointly benefit from public goods) depend on preferences and are therefore unequally distributed across couples. Couples where partners have a relatively strong preference for private goods can share fewer resources and need a higher level of household income so that both partners are not welfare-poor. I find that both of these aspects play a role. 46% of those who are welfare- but not income-poor live with a non-welfare-poor partner, highlighting the role of intra-household inequality. 54% of them live in

[5]For example, the Pareto weights in structural models are often determined by either market-clearing on a frictionless marriage market (e.g. Chiappori, Costa Dias, and Meghir (2018), Gayle and Shephard (2019) or Reynoso (2018)) or bargaining in a search model (e.g. Mazzocco, Ruiz, and Yamaguchi (2013) or Low et al. (2018)). In addition, Cherchye et al. (2017) and Cherchye et al. (2020) use marriage market conditions in a revealed preference setting in order to identify consumption allocations. Browning et al. (2021) use marriage market conditions to identify the unobserved values of 'match quality' in each couples.

households where both partners are welfare-poor, which suggests low economies of scale.

To quantify the role of preference heterogeneity, I also compute poverty based on the so-called *Empirical Sharing Rule* (ESR, Lise and Yamada (2019)). The ESR is defined as the sum of individual consumption and the market value of leisure and housework. This welfare measure allows to capture unequal sharing and differences in time use, but it still ignores preference heterogeneity as no information on preferences is used to compute it. If the result that income only identifies 61% of welfare-poor individuals was primarily driven by unequal sharing or time use, but not by preferences, the ESR should identify most welfare poor individuals, as it takes both of these factors into account. I find that the ESR identifies 81% of welfare-poor individuals. This is a substantial improvement over using household income and suggests that a significant part of the difference between income and welfare-poverty is driven by unequal sharing and time use. At the same time, it still leaves 19% of the welfare-poor unidentified. This highlights that using information on preferences is important to identify welfare-poor individuals.

To demonstrate the policy relevance of using individual-level welfare measures, I further study a hypothetical increase in the minimum wage and its effect on poverty among minimum wage earners. A well-known fact about minimum wages is that many minimum wage recipients live in households which are not income-poor, as they can live in couples with partners who have higher earnings. From the perspective of household income, this suggests that raising the minimum wage (MW) would not reduce poverty among MW earners much, as many of them are not income-poor in the first place. However, MW earners could still be poor in terms of their individual welfare in these cases, due to intra-household inequality and economies of scale as discussed above.

Using the estimated model, I simulate a 20% increase in the minimum wage and study the impact on welfare-poverty. The main result from this analysis is that the MW increase reduces welfare-poverty both for MW wage earners in less well-off families (where both partners have a low wage) and in families where the partner of the MW earner has a higher wage. For example, the reform reduces welfare-poverty rates by 6pp when the partner also earns the MW wage and by 7pp when the partner is higher up in the wage distribution and on average earns £22, which is well above the national average. These results imply that minimum wage increases benefit MW earners more strongly than one would expect based on household income, which increases their redistributive value in terms of lowering welfare-poverty rates.

Related Literature. This paper makes several contributions. First, the main contribution of the paper relates to the growing literature on measuring consumption and welfare on the level of individuals. Chiappori and Meghir (2015) and Chiappori and Meghir (2014) theoretically discuss individual welfare measures and propose the *Money-Metric Welfare Index* (MMWI) as a measure that accounts for intra-household inequality. On the empirical side, various methods for recovering individual consumption from the data have been proposed (e.g. Browning, Chiappori, and Lewbel (2013), Dunbar, Lewbel, and Pendakur (2013), Cherchye et al. (2017), Cherchye et al.

(2015), Lechene, Pendakur, and Wolf (2019), Cherchye et al. (2020)). Cherchye et al. (2018) propose a method for estimating the MMWI based on a model of consumption and leisure. The resulting estimates have often been used for an individual-level poverty analysis (e.g. Cherchye et al. (2018), Cherchye et al. (2015), Cherchye et al. (2020)). There has also been a sizeable literature on estimating individual consumption in developing countries (e.g. Bargain, Donni, and Kwenda (2014), Calvi (2020), Penglase (2021), Brown, Calvi, and Penglase (2021), Tommasi (2019)). These studies have largely focused on identifying individual consumption levels and define poverty based on consumption, rather than on welfare more broadly. The contribution of my paper is to jointly model consumption sharing, time allocation and preference heterogeneity and focus on welfare as measured by the MMWI. In particular, the role of preference heterogeneity for an individual-level welfare analysis has not been studied before and I find preference heterogeneity to be important. This paper also differs in terms of methodology, by using a more structural approach whereas other papers in this literature have used revealed preferences approaches or microeconometric estimation. The structural approach comes with different strengths, such as being able to conduct policy experiments, and thereby complements the existing literature.

Second, a recent literature has studied redistributive policies from the perspective of collective models. These papers have focused on the tax and transfer schedule or specific welfare policies such as the EITC or the 1996 US welfare reform (e.g. Gayle and Shephard (2019), Mazzocco, Ruiz, and Yamaguchi (2013), Low et al. (2018), Bronson and Mazzocco (2019), Obermeier (2019)). My paper contributes by studying the impact of minimum wage increases in a collective household model. There has been very little research on minimum wages from a family perspective. The only paper is Fields and Kanbur (2007), who theoretically analyse the impact of minimum wages on (consumption) poverty in the presence of exogenous income sharing. My paper differs by its focus on welfare poverty and the endogenous allocation of consumption and time.

Finally, the paper contributes to the broader literature on collective household models by building a model which can account for unobserved preference heterogeneity. Previous literature has allowed for limited preference heterogeneity which is typically tied to observables, such as gender, age or education (e.g. Lise and Yamada (2019), Cherchye, De Rock, and Vermeulen (2012)). While this set-up is well-suited for addressing certain questions, such as the long-standing question whether reallocation of resources to women changes household behaviour (e.g. Lundberg, Pollak, and Wales (1997)), it can restrictive in other applications as there is no heterogeneity conditional on observables. In other words, these models imply that all women care either more or less about public goods than all men (conditional on e.g. age and education). My model allows for preference variation within these observable cells, which is important for studying heterogeneity and distributional questions. Collective models have been applied in various contexts, such as tax and transfer policy (as discussed above), education policy (Chiappori, Costa Dias, and Meghir (2018)), divorce laws (Voena (2015), Reynoso (2018), Foerster (2020)) and macroeconomics (Doepke and Tertilt (2016), Knowles (2012)), and my paper highlights that unobserved preference heterogeneity could be included in such applications.

The rest of this paper is organised as follows. Section 2 describes the model, which is estimated in section 3. Section 4 shows the details of the policy simulations and discusses the results. Section 5 concludes.

2 Model

2.1 Overview

The purpose of the model is to study the link between household income and the welfare of the individuals living in the household. While household income can be measured directly, a model is required in order to define welfare. In addition, having a model is important to account for preference heterogeneity, as preferences are not directly observed in the data.

In line with standard setups in the collective model literature (e.g. Lise and Yamada (2019)), the model considers three particular goods: personal consumption (c_i), leisure (l_i) and a so-called 'home good' (D), which is produced using money and time and captures all joint activities of the couple (such as common expenditure and time spent working in the household). In each model period, households make a static choice between these goods.[6] This decision problem is embedded in a simple dynamic marriage market (building on e.g. Knowles (2012)).[7] This endogenises the allocation of decision power within the couple and determines the matching of the unobserved preference types.

This framework will make it possible to compare welfare across individuals based on the utility they receive from private consumption, leisure and the home good, i.e. $u_i(c_i, l_i, D)$ (converted into a monetary index), and compute poverty rates in both household income and welfare.

2.2 Demographics

The model contains a two-sided matching market with two populations, which can for example be thought of as men and women.[8] In each model period, a new cohort of these individuals ($g \in \{f, m\}$) is born. They start their lives as singles and they live until a terminal period T.[9] The

[6] Note that the model abstracts from savings. Savings are another reason why household income does not fully reflect welfare (Krueger and Perri (2006)), as some low-income households are affected by adverse wage shocks, but are able to maintain their living standard via savings. In a model with savings, the arguments from this paper about the link between household income and individual welfare would similarly apply to the link between household *expenditure* (i.e. $c_i + c_j + q$) and welfare.

[7] The model is most closely related to Obermeier (2019) and adds multi-dimensional preference heterogeneity while abstracting from asset accumulation. In addition, other related recent marriage market models are Ciscato (2019) and Shephard (2019).

[8] Note that the model assumes that individuals cannot search for partners within their own population. See Ciscato, Galichon, and Goussé (2020) for an equilibrium model of the same-sex marriage market.

[9] As the focus of the model is on static consumption and time allocation choices, the finite-horizon structure is not critical for the model. In practice, it turns out that solving a finite-horizon version of the model has computational advantages over an infinite-horizon model.

marriage market is an equilibrium random search market. In each period, singles meet someone from the distribution of available singles. Matching is restricted to the same age group, so that the age of partners is always identical. Couples exogenously have children ($b \in \{0,1\}$). Individuals discount the future at rate β.

Individuals differ in their multi-dimensional *type* ($t_i = (a_i, p_i)$), which consists of the *ability type* (a_i) and the *preference type* (p_i). Ability determines the wages that the individual can obtain in the labour market. While ability corresponds directly to wages in the data, the preference type p_i of an individual is unobserved in the data. Ability is drawn from a distribution F^A and the preference types is drawn according to F^P.

2.3 Goods and Preferences

Individuals derive utility from personal consumption (c), leisure (l) and the home good (D). Personal consumption and leisure are private goods and the home good is a public good, which means that once the household chooses to produce a level D, both partners consume D units of the home good. Individual differ in their preferences about the goods. The *preference type* of each individual is a vector $p_i = (\alpha_i^c, \alpha_i^l, \alpha_i^D)$ of preference coefficients for each good, with $0 < \alpha_i^K < 1$. The coefficient for private consumption (α_i^c) is normalised such that all coefficients add up to 1 (i.e. $\alpha_i^c = 1 - \alpha_i^D - \alpha_i^l$), so that there is effectively heterogeneity in a two-dimensional space (illustrated in Figure 1).[10]

FIGURE 1: Illustration of preference space

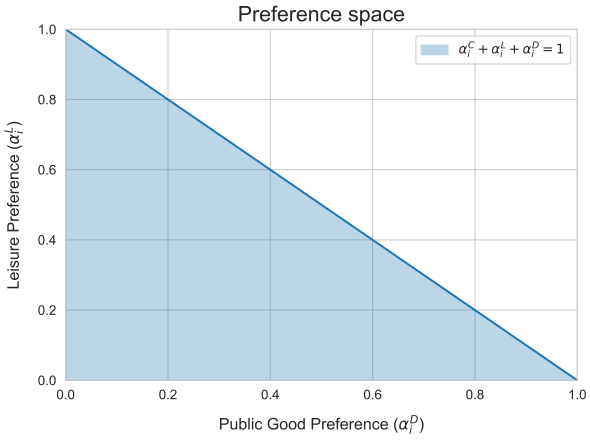

Notes: The figure shows the feasible combinations of preference parameters. The preference for private consumption is implicitly defined as $\alpha_i^C = 1 - \alpha_i^D - \alpha_i^L$ and all preference coefficients must be in the open interval $(0,1)$.

The preference coefficients determine the weight of each good in the utility function, where γ is

[10] This can be thought of as a *no-utility-monster restriction*, as it rules out that individuals obtain more utility from each good.

a common curvature parameter:

$$u_i(c_i, l_i, D) = \alpha_i^c \frac{c_i^{1-\gamma}}{1-\gamma} + \alpha_i^l \frac{l_i^{1-\gamma}}{1-\gamma} + \alpha^D \frac{D^{1-\gamma}}{1-\gamma}$$

Here, the notation makes clear that c_i and l_i are private and D, the home good, is determined on the household level, so that all members of household h consume the same amount of the home good. In this sense, the home good is public (i.e. non-rival and non-excludable among family members). The *utility* that a person obtains from the public good may be fairly different for each partner, as this also takes the individual preference weight into account ($\alpha_i^D \frac{D^{1-\gamma}}{1-\gamma}$). For example, an extreme case would be the situation where one partner does not care at all about the public good $\alpha_i^D = 0$. In this case, the home good would not contribute to this person's utility, even though it is a public good.

The production technology for the home good takes both *aggregate* domestic work $\bar{d}$ and public expenditure q as inputs:

$$D = \bar{d}^{\omega_H} q^{1-\omega_H}$$

For singles, total domestic work is equal to their individual domestic hours. For couples, the domestic hours of each spouse are aggregated with another CES function:

$$\bar{d} = (\theta_f(b) d_f^z + \theta_m(b) d_m^z)^{\frac{1}{z}}$$

The home production technology depends on the presence of children (b), which captures that the presence of children affects the division of labour within the household. The model includes gender differences in domestic hours through gender-specific 'productivity' coefficients $\theta_f(b)$. Of course, this does not necessarily reflect actual differences in productivity, but could also represent the presence of gender norms, which make women's time valued more strongly for home production.[11] In addition to the gender-specific productivity coefficients, having kids is also allowed to shift the preference parameter for domestic hours with a scale parameter η_b, which captures that couples with kids increase the amount of resources spent on the public good.

2.4 The Decision Problem of Singles

Singles solve a static decision problem on expenditure on personal and public consumption and time use in each period of singlehood. They also face the dynamic decision of whether to start a relationship or keep searching when they meet other singles, as will be explained in more detail

[11] An alternative specification of the gender norm would be to introduce a utility penalty for using the male home production input, reflecting the social disutility from this choice.

in section 2.6. The decision problem of singles is given by:

$$\begin{aligned} \max_{c,q,l,d,h} \; & u_i(c,l,D) \\ c+q &= w_i h \\ l+d+h &= 1 \\ D &= D(q,d) \end{aligned}$$

The time budget of the individual is normalised to 1. The solution to this problem leads to the indirect utility function $U_i^S(\omega_i^S)$, where ω_i^S is the state vector of singles, which contains gender, ability, preference type and whether the single has children from a previous relationship:

$$\omega_{it}^S = (g_i, a_i, p_i, b_{it})$$

As singles meet potential partners and can form couples, the value of singlehood at age (t) takes into account both the solution of the choice problem and the expected future value (EV_{t+1}^S) which also includes future marriage market outcomes (see section 2.6):

$$V_{it}^S(\omega_{it}^S) = U_i(\omega_i^S) + \beta \text{EV}_{t+1}^S(\omega_{it}^S)$$

2.5 The Decision Problem of Couples

Couples similarly solve a static choice problem about expenditure and the time allocation of each partner. Suppose that woman i is in a relationship with man j, which is denoted as household $h = \{i, j\}$. Households are characterised by the ability and preference type of each partner, the love shock of the couple (θ_{ht}), the Pareto weight of the woman λ_{ht} and the presence of children:

$$\omega_{ht}^C = (a_i, a_j, p_i, p_j, \theta_{ht}, \lambda_{ht}, b_{ht})$$

The couple maximises a weighted sum of utility, where the weight of each partner is determined by the Pareto weight λ_{ht}:

$$\begin{aligned} \max_{c_i,c_j,q,l_i,l_j,d_i,d_j,h_i,h_j} \; & \lambda_{ht} u_i(c_i,l_i,D) + (1-\lambda_{ht}) u_j(c_j,l_j,D) \\ c_i + c_j + q &= w_i h_i + w_j h_j \\ l_i + h_i + d_i &= 1 \\ l_j + h_j + d_j &= 1 \\ D &= D(q, d_i, d_j) \end{aligned}$$

Note that the utility functions of each partner (u_i and u_j) can differ because of preference heterogeneity. The Pareto weight can be interpreted as 'bargaining power' within the household.

For example, if the woman has a high Pareto weight, the household will allocate a lot of private consumption and leisure to her. In addition, a high value for the Pareto weight gives her a lot of decision power over the public good, as she may value the public good differently from her partner. If the woman cares less about the home good than the man, an increase in λ_{it} will decrease money and time spent on the producing the public good.[12] Figure 2 illustrates these effects. In this example, the parameters are set such that the woman has a lower preference for the public good. As the Pareto weight λ_{it} increases, her share of private consumption increases while the fraction of expenditure devoted to the public good decreases.

FIGURE 2: Illustration of couple choice problem

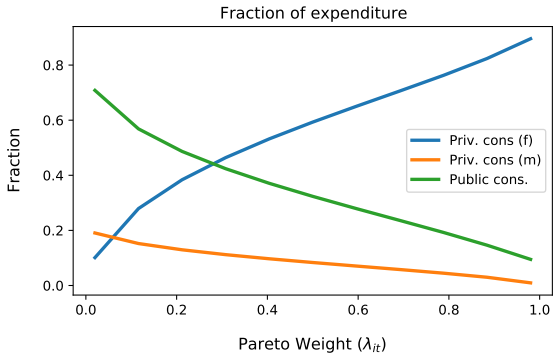

Notes: This figure gives an example how the share of budget devoted to private consumption of each partner and to the household-level public good varies with the Pareto λ_{it}. In the example, the female partner has a lower preference for the public good.

Analogously to singles, the utility obtained from the choice problem is denoted as $U_i^C(\omega_{ht}^C)$ (h being the household index). For each partner $p \in \{i, j\}$, the value of being in a couple further contains love (Θ_{ht}) and the continuation value.

$$V_{pt}^C(\omega_{ht}) = U_p^C(\omega_{ht}^C) + \Theta_{ht} + \beta \text{EV}_{p,t+1}^C(\omega_{ht}^C)$$

The continuation value is defined more precisely in section 2.6. Love has a stochastic component (θ_{ht}), which is drawn at the time of marriage and then evolves according to a random walk process. In addition, there is also *homophily* term in labour market ability, which allows the model to capture assortative matching in terms of ability:

$$\Theta_{ht} = \theta_{ht} + \frac{\kappa}{1 + (a_i - a_j)^2}$$

The sorting term leads to additional utility κ if the ability types are equal and this extra utility decreases the further the two types are apart. The interpretation of the term could for instance be that higher ability correlates with other traits or education which are valued on the marriage

[12] See Blundell, Chiappori, and Meghir (2005) for a thereotical discussion of the conditions required for this effect.

market and lead to sorting.[13]

2.6 The Marriage Market

2.6.1 Meetings and Matching

In the beginning of each period, singles are randomly matched with another single from their own age group. The distribution of available singles at each age corresponds to those individuals who stayed single in previous periods (or already separated). These distributions are equilibrium objects (see section 2.6.3).

Suppose that single i meets single j. The potential couple draws a love shock θ from a normal distribution ($\theta \sim N(\mu_l, \sigma_l)$). The love shock captures the non-economic quality of the relationship. It provides a motive for searching for a partner with a high realisation of the love shock, as it is newly drawn for each potential partner that an individual meets.

Upon meeting, singles observe the value of the love shock and all characteristics of each partner (i.e. ω^S_{pt} for each person $p \in \{i, j\}$). They can decide whether they want to form a couple and on the value of the Pareto weight. If they get together, their joint state space is ω^C_{ht} (which includes the Pareto weight) and each partner p obtains utility $V^C_{pt}(\omega^C_{ht})$, where V^C_{pt} is determined by the decision problem of couples from the last section. For a relationship to be viable, the value of being in the relationship must be higher than the value of singlehood for partner p:

$$V^C_{pt}(\omega^C_{ht}) \geq V^S_{pt}(\omega^S_{pt})$$

If there is no Pareto weight λ_{ht} that ensures that both singles prefer the relationship, they remain single. Otherwise, the initial Pareto weight is determined by Nash bargaining. The bargaining problem is described by the following maximisation problem (note that the Pareto weight λ is included in ω^C_{ht}):

$$\tilde{\lambda} = \text{argmax}_\lambda \left(V^C_{it}(\omega^C_{ht}) - V^S_{it}(\omega^S_{it}) \right) \cdot \left(V^C_{jt}(\omega^C_{ht}) - V^S_{jt}(\omega^S_{jt}) \right)$$

After the relationship has started, the Pareto weight stays constant over time unless one partner wants to leave the relationship (see below). This is the standard limited commitment structure as described in more detail e.g. in Chiappori and Mazzocco (2017).[14]

Given these marriage market decisions, the continuation value of singlehood can be expressed

[13] In practice, κ is identified as the residual required to match the extent of wage sorting observed in the data.

[14] Assuming limited commitment (relative to no commitment) increases the scope for transfers via the Pareto weight, as one person can promise the other one a high level of utility in all periods. As a result, the extent of commitment affects the matching patterns that arise in equilibrium.

more explicitly as the expectation over the marriage market outcomes in the next period:

$$\begin{aligned}\mathrm{EV}_{it}^{S}(\omega_{it}^{S}) = \int_{\omega_{j,t+1}^{S},\theta_{h,t+1},b_{t+1}} & M(\omega_{h,t+1}^{C}) \cdot V_{i,t+1}^{C}(\omega_{h,t+1}^{C}) \\ + & (1 - M(\omega_{h,t+1}^{C})) \cdot V_{i,t+1}^{S}(\omega_{it}^{S}) \, \mathrm{d}F(\omega_{j,t+1}^{S}, \theta_{h,t+1}, b_{t+1})\end{aligned}$$

In words, the expected future utility is the integral over all potential partners that individuals might meet in the next period, the love shock, and whether the individuals' children grow up (if they have children). For each of these cases, the individual obtains utility $V_{i,t+1}^{C}(\omega_{h,t+1}^{C})$ if the meeting results in a relationship and utility $V_{i,t+1}^{S}(\omega_{i,t+1})$ otherwise.

2.6.2 Separation

Couples can also terminate their relationship if the love shock changes. Partner $p \in \{i, j\}$ can unilaterally initiate a separation if their value from singlehood is larger than the value of staying in the relationship:

$$V_{pt}^{C}(\omega_{ht}^{C}) < V_{pt}^{S}(\omega_{pt}^{S})$$

Following the standard limited commitment setup, couples are allowed to renegotiate on the weight. There are three cases. First, if both partners prefer separation, they separate. Second, it can happen that for a given value of the love shock, only one partner wants to separate (because of heterogeneity in wages and preferences).[15] If it is possible to adjust the Pareto weight to make them indifferent between leaving and staying, the couple renegotiates and stays together. Third, if it is not possible to find a new Pareto weight such that both partners want to stay, the couple separates.

The expected continuation value from being in a couple is given by:

$$\begin{aligned}\mathrm{EV}_{pt}^{C}(\omega_{ht}^{C}) = \int_{\theta_{h,t+1},b_{t+1}} & (1 - S(\omega_{h,t+1}^{C})) \cdot V_{p,t+1}^{C}(\omega_{h,t+1}^{C}) \\ + & S(\omega_{h,t+1}^{C}) \cdot V_{p,t+1}^{S}(\omega_{p,t+1}^{S}) \, \mathrm{d}G(\theta_{h,t+1}, b_{t+1})\end{aligned}$$

Note that the state space of the couple in the next period ($\omega_{h,t+1}^{C}$) can differ from the current period when a renegotiation takes place. Thus, the expected future value for couples takes the possible of separation ($S(\omega_{h,t+1}^{C}) = 1$) as well as changes in the Pareto weight into account.

[15] For example, the lower wage partner is more likely to want to stay in the relationship for a given value of love, as they have a lower living standard when single. Similarly, people who have a strong preference for the public good value being in a relationship more.

2.6.3 Marriage Market Equilibrium

The marriage market equilibrium is a rational expectation equilibrium. The basic equilibrium requirement is that the expected distributions of singles that agents take into account when making decisions is equal to the actual distributions which emerge from these decisions. In order to solve the life-cycle problem, a guess for the distribution of singles at each age group is needed, as agents need to take the probabilities of spouses they might meet in the future into account. Conditional on this guess, the actual distribution of singles can be computed, which in turn feeds back into the life-cycle problem. As a result, the equilibrium can only be solved for numerically by fixed-point iteration. The numerical solution of the model is described in more detail in the appendix.

Equilibrium Definition. A stationary equilibrium consists of distributions of singles, policy functions for singles and couples and matching rules such that

1. the policy functions $(c, q, h, l, d) = P^S(\omega^S)$ solve the problem of singles
2. the policy functions $(c_i, c_j, q, h_i, h_j, l_i, l_j, d_i, d_j) = P^M(\omega_h^C)$ solve the problem of married couples
3. separation and rebargaining $(D, \tilde{\lambda})$ occur according to the limited commitment procedure
4. the matching rule (m, λ) satisfies the participation constraints and the bargaining solution, where m is an indicator for starting a relationship and λ the initial Pareto weight
5. the implied distributions of singles for each gender and age, $\Lambda_{t,g}(\omega^S)$ $(g \in \{f, m\}))$, are consistent with the distributions that are used to determine the optimal choices and value functions from (1) - (4)

2.7 Welfare Measurement with Preference Heterogeneity

2.7.1 Definition of Money-Metric Welfare Index (MMWI)

Before turning to the quantitative part, it is useful to discuss how to think about individual welfare in the context of the model. Measuring welfare in the context of preference heterogeneity leads to conceptual questions about how to compare utility levels across individuals. Chiappori and Meghir (2015) have recently proposed the so-called *Money-Metric Welfare Index* (MMWI) for welfare comparisons within the family. The idea is to compute the hypothetical resources, in terms of the expenditure function, which an individual would need as a single in order to achieve the same utility level as in marriage. This then leads to a monetary welfare measure which can easily be compared across individuals.

Formally, suppose an individual reaches the utility level $\bar{u} = u_i(c, l, D)$ while living in a couple. Note that $\bar{u}$ does not include the value of the love shock, in order to focus on economic inequality. The MMWI is the *full* income (i.e. consumption plus the cost of time use) which the individual would need as a single in order to obtain utility level $\bar{u}$. As the time budget is normalised to 1, full income is given by the wage rate of the individual (w). In order to compute the MMWI, one first needs to solve for the indirect utility function of the single problem:

$$\begin{aligned} V_i(w) &= \max_{c,q,l,d} u_i(c, l, D) \\ c + q + wl + wd &= w \\ D &= D(q, d) \end{aligned}$$

Note that u_i depends on the preferences of the individual. Then, the MMWI is implicitly defined by:

$$V_i(\text{MMWI}_i) = \bar{u}$$

In the analysis, the MMWI will be expressed as the hourly wage rate of the individual. It could equivalently be expressed as full income by multiplying it with the number of total available hours.

The interpretation of the MMWI is that individuals are compared according to their full income (i.e. their wage). According to the MMWI, those individuals with the lowest wages are the least well-off in the population, which captures inequality in the budget constraint which individuals are facing. For individuals who live in couples, the MMWI adds their share of the surplus onto their wage. For example, a coupled individual whose hourly wage is £10 might have an MMWI of £13 if being in the relationship makes them 30% better off (due to economies of scale and intra-household sharing). In other words, this individual is as well off as being single and earning a wage of £13. Some examples for particular parameter values of the model are discussed in section 2.7.2.

Note that there are multiple ways of comparing welfare in the presence of preference heterogeneity. The literature in welfare economics has introduced several welfare measures which capture different ethical judgments about taking work vs non-work preferences into account (see e.g. Fleurbaey (2006), Decoster and Haan (2015)). The MMWI is similar to one of the options proposed by Fleurbaey (2006), the "Wage Rate Criterion". I discuss the other criteria from Fleurbaey (2006) and the underlying ethical judgments in more detail in appendix A and show how the main results are affected by the choice of the normative framework. In the main text, I focus on the MMWI as this makes the analysis more comparable to previous work on intra-household inequality (Cherchye et al. (2018), Chiappori and Meghir (2015)).

2.7.2 Examples

To get a better intuition about the MMWI and welfare comparisons in the presence of preference heterogeneity, it is useful to discuss a few numerical examples and look at the allocations and values of the MMWI in each case.

Table 1 shows numerical examples for different couples. In each couple, the woman and the man have different preference types, which belong to a set of three possible types. The first type is 'consumption-oriented' and places a higher weight on consumption and a lower weight on the other two goods. The second and third type are 'leisure-oriented' and 'home-oriented'. The corresponding values of these preference types (consisting of α^c, α_l and α_D) are given by:

$$\begin{aligned} p_1 &= (0.45, 0.275, 0.275) \\ p_2 &= (0.275, 0.45, 0.275) \\ p_3 &= (0.275, 0.275, 0.45) \end{aligned}$$

For example, the first column shows the case where both the man and the woman are consumption-oriented. Men have an hourly wage of £10 and women's wage is set to £8. The Pareto weights are determined by Nash bargaining, where the outside option is the value each person would obtain in singlehood. In columns (1)-(3), the table first shows the outcomes for three assortative matches (i.e. where both partners have the same preference types). In these cases, the woman consumes less and has less leisure than the man. For example, for the consumption-oriented couple, the woman consumes 19% less and has 7% less leisure. The MMWI aggregates these differences into a single welfare statistic, also taking the utility from the home good into account (which is equal for both partners in this case, as they have similar preferences). According to the MMWI, the woman is 18% less well off than the man in this case. Columns (4)-(6) show the case of mixed couples. These cases are interesting as it is not clear based on observing consumption or leisure alone who is better off. For example, in the case of a consumption-oriented woman and a leisure-oriented man, the woman consumes more but has less leisure. In all of these mixed cases, the MMWI concludes that the woman is worse off than her partner. This results from the gender wage gap, as women's wages are £2 lower in this example. For comparison, the table also shows the underlying (non-money-metric) utility levels for each partner.

Also note that the surplus from the relationship differs across couples with different preference types. Measuring the surplus (i.e. economies of scale) has been an important topic in the collective model literature (see e.g. Cherchye et al. (2018)). In the context of poverty, economies of scale are important as households with high economies of scale can lift both individuals above the poverty threshold with fewer resources. The model in this paper allows to capture that economies of scale depend on preferences and vary accordingly in the population. In particular, the surplus is higher when when both partners are home-oriented (column 3), as this allows them to share more resources.

Finally, the table also shows the values for the *Rente Criterion* from Fleurbaey (2006), which is an alternative way of comparing welfare in the presence of preference heterogeneity. Focusing on the last two rows of the table illustrates that the Rente Criterion and the MMWI can differ in their welfare conclusions about who in the couple is better off. While both criteria are give similar results when preferences are identical (rows 1-3), they can differ when preferences are heterogeneous (rows 4-6). These issues, and the implications of the choice between different welfare criteria for the rest of the analysis, are discussed in more detail in appendix A.

3 Estimation

3.1 Data

To obtain information on the time use of singles and couples, I use data from the *United Kingdom Time Use Survey* (UKTUS). The data contains detailed time diaries where individuals record their activities during the day. Individuals are surveyed on two days (a weekday and a weekend day). The data contains some demographic background variables, including the number of children. To increase sample size, I pool the two waves of 2014-2015 and 2000-2001. Table 2 shows the summary statistics for the sample. In addition, Figure 3 shows the cross-sectional variability in time use choices of singles.[16]

TABLE 2: Summary Statistics (UKTUS)

	Single Men	Single Women	Coupled Men	Coupled Women
Work hours	28.93	22.08	39.03	24.01
Leisure Hours	42.80	35.48	30.11	30.95
Home Hours	12.65	24.00	12.70	27.94
Children in HH	0.07	0.44	0.55	0.55
N	703.00	1163.00	3227.00	3227.00

Notes: The table shows the summary statistics for the UKTUS sample. Time use is reported in weekly hours.

Since there is no information on individual earnings or wages in the time use survey, I also use *The UK Household Longitudinal Survey* (USOC). As USOC is a panel, it allows to compute moments

[16] An interesting extension of the analysis would be to also look at the variability in consumption behaviour. The Family Resources Survey (FES) contains very detailed consumption data for the UK and it is possible to classify expenditure into public and private expenditure following Lise and Seitz (2011). This would allow to analyse heterogeneity in the ratio of public and private consumption goods among singles as an additional indicator of preference heterogeneity for the public good.

TABLE 1: Welfare Measurement - Examples (Couples)

Preferred Goods (f / m)	(1) c / c	(2) l / l	(3) D / D	(4) c / l	(5) c / D	(6) l / D
(a) Consumption						
Income	4590.2	3653.9	4096.7	4063.8	4298.8	3916.7
c_f	1489.3	1077.6	1058.7	1489.6	1387.4	1004.2
c_m	1849.0	1325.2	1318.6	1325.3	1369.0	1368.3
C	1252.0	1251.2	1719.3	1249.0	1542.4	1544.1
(b) Time Use						
h_f	54.8	42.5	45.4	54.8	50.8	39.4
h_m	66.3	53.7	62.0	53.7	62.5	62.5
l_f	31.0	43.3	30.6	31.0	28.9	40.3
l_m	33.2	45.9	32.9	45.9	34.1	34.1
d_f	26.2	26.2	36.0	26.1	32.3	32.3
d_m	12.5	12.5	17.1	12.5	15.4	15.4
(c) Welfare						
λ	0.42	0.42	0.42	0.42	0.38	0.39
$MMWI_f$	9.33	9.7	10.18	9.33	9.51	9.96
$MMWI_m$	11.36	11.76	12.1	11.75	11.42	11.42
u_f	-3.67	-3.65	-3.71	-3.68	-3.66	-3.63
u_m	-3.48	-3.5	-3.56	-3.51	-3.61	-3.61
R_f	1370.45	1044.82	1240.78	1369.41	1392.58	1068.16
R_m	1621.22	1227.52	1442.23	1226.57	1370.69	1371.13
Surplus ($MMWI$), in pp	14.97	19.25	23.83	17.1	16.28	18.81
(d) Relative Welfare						
Relative c	-0.19	-0.19	-0.2	0.12	0.01	-0.27
Relative l	-0.07	-0.06	-0.07	-0.32	-0.15	0.18
Relative $MMWI$	-0.18	-0.18	-0.16	-0.21	-0.17	-0.13
Relative R	-0.15	-0.15	-0.14	0.12	0.02	-0.22

Notes: The table shows examples for the welfare measurement for couples with different combinations of preferences. λ is the Pareto weight of the couple. h_g, l_g and d_g are market hours, leisure and home hours. Time use is reported in weekly hours, the MMWI in £ per hour and income, consumption and the Rente Criterion (see Appendix A) in £ per month. Relative variables are computed as e.g. $\frac{c_f - c_m}{c_m}$ in the case of private consumption and reported in %.

FIGURE 3: Variability of Time Use - Singles

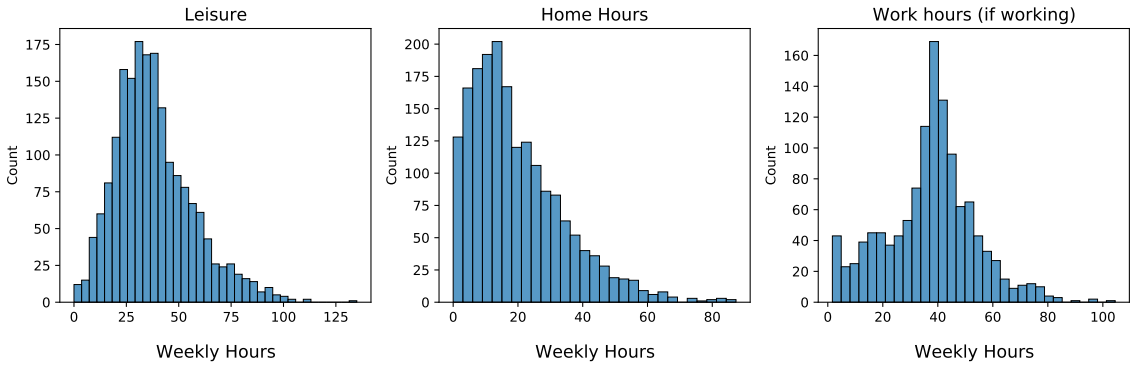

Notes: The histograms show time use choices of singles in the UKTUS data.

on separation and marriage rates.

3.2 Model specification

In practice, the time horizon is set to $T = 30$. As a result, individuals start their lives at age 20 and each period correspond to two years. To estimate the model, one further needs an empirical specification for the preference distributions. Recall that an individuals' preference type is a vector $(\alpha_i^C, \alpha_i^L, \alpha_i^D)$, where one coefficient can be normalised so that $\alpha_i^D = 1 - \alpha_i^L - \alpha_i^C$. As illustrated by Figure 1, this amounts to specifying a probability distribution over a triangular space. In practice, I draw auxiliary normal random variables χ^K ($K \in \{C, L, D\}$) with mean μ_K and variance σ_K. μ_D is normalised to 1. Then the preference coefficients are defined as:

$$\alpha^K = \frac{\chi^K}{\sum_k \chi^k}$$

This leads to a continuous distribution of preference types. The final step is to discretise this distribution using N_P points. Unlike for univariate distributions, there is no standard procedure for discretising a multivariate distribution. As points, I choose a linear grid for the first preference dimension and add the second dimension by computing the conditional quantiles conditional on each point in the linear grid. Finally, the probabilities of each of the points are computed by assigning each point from the continuous distribution to the closest point of the 2D grid. Applying this procedures allows to replicate well the first and second moments of the continuous distribution. Taken together, this procedure leads to a preference distribution with 5 parameters $(\mu^C, \mu^L, \sigma^C, \sigma^L, \sigma^D)$.

TABLE 3: Estimated Parameters

Parameter	Description	Value	Target moment
μ^L	Preferences, mean parameter 1	0.44	Mean work/leisure/home hours in couples
μ^C	Preferences, mean parameter 2	0.32	Mean work/leisure/home hours in couples
σ^C	Preferences, variance parameter 1	0.61	Variabilility of singles choices
σ^L	Preferences, variance parameter 2	0.00	Variabilility of singles choices
σ^D	Preferences, variance parameter 3	0.36	Variabilility of singles choices
η_b	Preference scale for home good, with kids	0.26	Mean work/leisure/home hours in couples
$\theta(b=0)$	Home productivity, women, no kids	0.88	Relative time use in couples
$\theta(b=1)$	Home productivity, women, with kids	0.94	Relative time use in couples
ω_H	Weight on time, home production	0.53	Relative time use in couples
μ_L	Mean, initial love shock	0.94	Marriage and separation rate
σ_l	Variance, love shock	0.93	Marriage and separation rate
κ	Homophily parameter	0.73	Within-HH variance in wages

Notes: This table shows the parameters which are jointly estimated to match the data moments and their values.

3.3 Estimation strategy and estimated parameters

Some model parameters are set externally. The yearly discount rate β is to 0.98 to capture standard practice in the literature. In addition, the curvature of the utility function (γ) is set to 1.5. The substitutability between home hours in the household production function (z) is set to -2, which is found to lead to a realistic relationship between relative wages and relative work hours.[17]

The remaining parameters are estimated to match a set of data targets, by minimising the distance between model and data moments. These parameters are listed in Table 3. Table 4 shows the target moments which are used to pin down the parameters.

The first set of parameters and moments relates to the marriage market. The parameters are the mean of the love shock (μ_L), the variance of the love shock (σ_L) and the homophily parameter (κ). The first two are estimated in order to match the fraction of individuals living in couples and the separation rate. The homophily parameter is estimated such that the fraction of total

[17]Note that this implies that home hours are complements in the household production function. This is required to match the cross-sectional relationship between relative work hours and relatives wages, as otherwise the home production function implies that women whose partner earns more than them work counterfactually low hours. More formally, this can be seen by deriving the first-order condition for the optimal ratio of home hours and solving for this ratio:

$$\frac{d_i}{d_j} = \left(\frac{w_i}{w_j}\frac{\theta_j}{\theta_i}\right)^{\frac{1}{z-1}}$$

In this equation, z directly governs the steepness of relative home hours with respect to wages. Only negative values of z can generate a reasonable cross-sectional profile of home hours.

wage inequality which is due to within-household inequality matches the data. Intuitively, if the homophily parameter takes a high value, matching will be very assortative and there will be little wage inequality within households (and vice versa). The remaining parameters relate to the home production and preference parameters. These are the 'productivity' parameters for women, depending on the presence of kids, the preference scale for the presence of kids, and the weight on time in the utility function. In addition, there are the means and variances of the preference distribution. All of these parameters are jointly estimated to match moments on work and home production hours of coupled men and women, separately by the presence of kids. Finally, the variance parameters of the preference distributions target the variability of time use choices of singles, where the mean absolute deviation (MAD) is chosen as an outlier-robust measure of variability. This captures the intuition that the variance among singles is informative about the extent of preference heterogeneity in the population. Note that the parameters of the preference distribution are auxiliary parameters and that the variance parameters cannot directly be interpreted as the variances of the resulting preference distributions (as preferences are computed as $\alpha^K = \frac{u^K}{\sum_k u^k}$, as discussed in the last section). The resulting preference distribution is best illustrated by Figure 4.

TABLE 4: Model Fit

	Data	Model
Work hours, women, mean, couples, no kids	32.08	26.32
Home hours, women, mean, couples, no kids	21.83	22.98
Work hours, women, mean, couples, with kids	22.86	24.11
Home hours, women, mean, couples, with kids	37.35	37.96
Work hours, men, mean, couples, no kids	38.31	41.88
Home hours, men, mean, couples, no kids	10.49	11.08
Work hours, men, mean, couples, with kids	40.33	50.96
Home hours, men, mean, couples, with kids	14.32	13.80
Work hours, MAD, Women, single	11.00	8.89
Leisure hours, MAD, Women, single	8.00	10.90
Home hours, MAD, Women, single	8.50	6.68
Marriage rate	0.78	0.83
Separation rate	0.04	0.06
Share within wage var	0.38	0.38

Notes: The table shows the fit of the model by comparing the targeted data moments to the corresponding model moments. "MAD" refers to Mean Absolute Deviation which measures the variability of time use choices.

3.4 Estimation Results: Assortative Matching and Selection into Relationships

Before turning to the welfare analysis, it is worth discussing the marriage market patterns, as these shed some more light on the model mechanisms and estimation. Table 5 describes the extent of assortative matching on preferences based on the estimated model by showing the rank correlations of spouses for each preference dimension. Focusing on the diagonal, there is a positive correlation for each preference type, which is strongest for the home good. This means that there is some assortative matching based on preferences, in the sense that for example individuals with a high preference for public goods in equilibrium are more likely to match with others who also like the public good. Overall, these patterns are best summarised as positive, but imperfect assortative matching. Figure 4 shows the joint distribution of preferences more fully.

TABLE 5: Assortative matching on preferences

Preferences (m) Preferences (f)	Consumption (α^C)	Leisure (α^L)	Home good (α^D)
Consumption (α^C)	0.05	0.07	-0.11
Leisure (α^L)	0.07	0.09	-0.14
Home good (α^D)	-0.10	-0.14	0.21

Notes: The table shows the rank correlations between the preferences of partners in the estimated model. The rows refer to preferences of the female partner and the columns to those of the male partner.

This is important for the rest of the analysis as assortative matching on preferences is an important determinant of intra-family preference differences. Under random matching, there would be as much variability in preferences between spouses as in the general population. If matching was perfectly assortative on preferences, there would be no within-couple differences even though preferences vary in the population. Table 6 illustrates this point more formally by showing the results from a variance decomposition of each of the preference coefficients.

FIGURE 4: Preference distributions in the estimated model

(a) Public good (matching)

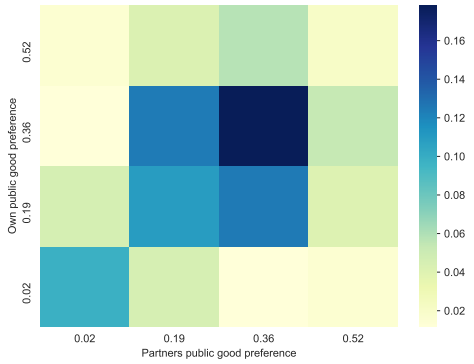

(b) Public good (coupled vs single)

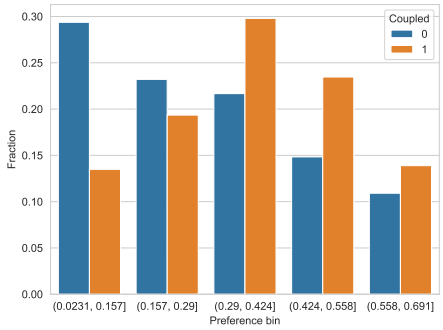

(c) Leisure (matching)

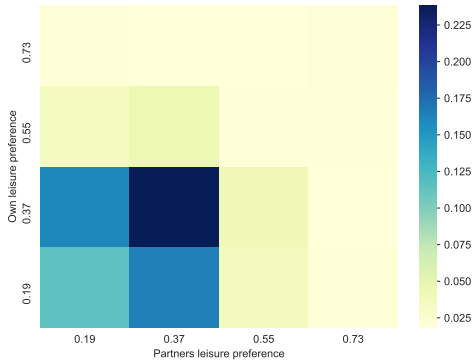

(d) Leisure (coupled vs single)

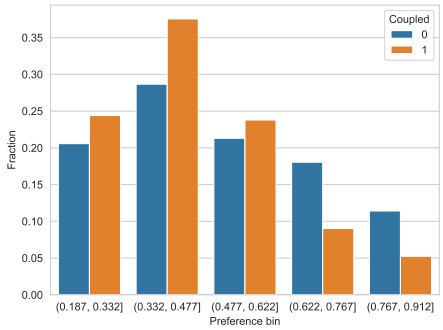

(e) Private consumption (matching)

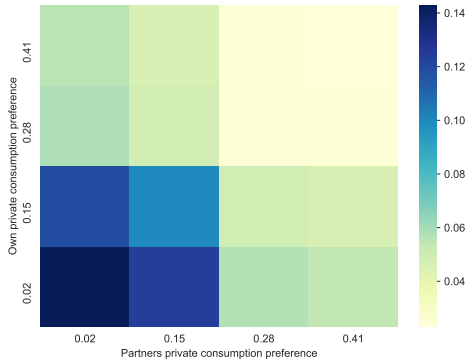

(f) Priv. cons. (coupled vs single)

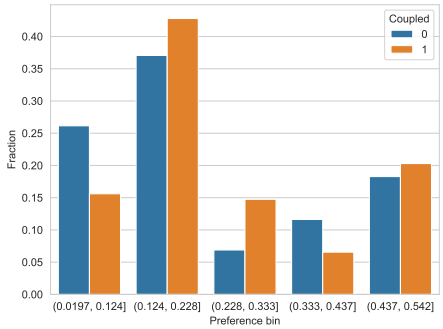

Notes: Each panel labeled as *matching* shows the joint distribution of preferences for the given preference dimension based on the estimated model. Both own and partners' preference coefficients are grouped into quartiles (as the underlying preference types are discrete) and the colour indicates the fraction of couples that have this combination of preference types. The panels labeled as *coupled vs single* show the distribution of preference types for all coupled/single individuals, again grouped into bins.

TABLE 6: Preference variability within and between households

	Within households	Between households
Consumption (α^C)	0.49	0.51
Leisure (α^L)	0.43	0.57
Home good (α^D)	0.30	0.70

Notes: The table shows the fraction of the variance of the preference coefficients ($\text{Var}(\alpha_i^K)$) that is within/between families, using the standard formula for variance decompositions. With positive assortative matching on preferences, there will be less variability within than between households..

It is worth highlighting that these patterns of assortative matching are identified through a combination of data and theory. On the hand, the data moments used in the estimation determine the preference distribution and parameters of the home production function. On the other hand, the marriage market model determines who gets married to whom based on the implied surplus from marriage. Only by combining these two steps, it is possible to pin down the joint distribution of preferences of partners, as there are no direct empirical measures of preferences.

Table 7 shows selection into relationships, as the incentives to get married differ for across preference types. The table illustrates that individuals who value the home good strongly are more likely to be in a couple, as the returns from being able to share public goods are higher for them. Conversely, those who value private goods (consumption and leisure) more strongly are relatively less likely to be in a relationship. Comparing the mean value of the preference coefficients, this results in married individuals having a slightly higher mean of the public good preferences and a lower value for leisure and private consumption preferences. Figure 4 shows the distribution of each of the preference coefficients for married and single individuals. The largest differences are seen for the public good preference. Among singles, there is substantially more mass at the lower end of the distribution whereas married individuals are more centered around the middle.

Another way of looking at tables 5 and 7 is asking to what extent the distribution of preferences among singles is informative about the distribution of preferences of couples. For example, if one had estimated a preference distribution for singles using data on singles only, what would this imply for the distribution of couples? This discussion highlights that there are two questions that need to be addressed for such an inference. The first is that the preference distribution of coupled individuals may be systematically different from the preference distribution among singles. This selection process is modeled explicitly through the marriage market. Table 7 shows that on average, the preferences of coupled and single individuals are fairly close in the estimated model. The second is that simply observing preferences in a sample of singles is not informative

TABLE 7: Selection into Relationships

	Correlation with being in couple	Mean (Singles)	Mean (Couples)
Consumption preference (α^C)	-0.04	0.204	0.185
Leisure preference (α^L)	-0.12	0.505	0.452
Home good preference (α^D)	0.15	0.291	0.363

Notes: The table reports the correlation between being in a couple and each of the preference dimensions based on the estimated model.

about the matching that occurs in couples. Even if there is a lot of preference heterogeneity among singles, there could be little within-household differences in preferences if matching is strongly assortative on preferences. As a result, a marriage market model is needed to map the preference distribution of singles into the joint distribution in couples.

4 Results

4.1 Welfare Analysis

4.1.1 Comparing Income and Welfare Poverty

In this section, I turn to the welfare analysis. The estimated model allows to measure welfare on the level of the individual by computing the Money-Metric Welfare Index (MMWI), which I use for an individual-level poverty analysis.[18] As the focus of the analysis is on intra-household discrepancies, the analysis is conducted using a simulated sample of couples based on the estimated model.

An individual is defined as income-poor if their household income is in the bottom 20% of the distribution. Note that income-poverty will always refer to household income instead of individual income and that it is computed using labour income rather than full income.[19] Income-poverty is effectively defined on the household level, as one partner being income-poor implies that the other partner is income-poor as well. Welfare-poor individuals are defined as those who are in the bottom 20% of the distribution of the MMWI. This holds the quantity of poor individuals constant across welfare measures, but the measures will have different implications for *who* is considered as poor. In contrast to income-poverty, welfare-poverty is defined on the level of the individual,

[18] To complement these results which focus on poverty (i.e. the lower end of the distribution), Tables 21 and 22 also show results for the whole distributions by focusing on rank correlations.

[19] As families are characterised by at least some degree of sharing, individual income often does not reflect welfare. In some related papers (e.g. Cherchye et al. (2018)), individuals are grouped into income bins according to full income (income if work hours are set to the maximum available time). Here, I use the actual labour income of the household, so that the results can be related more directly to the broader literature on inequality which compares labour income across families.

which means that being welfare-poor does not imply that one's partner is also welfare-poor (this will be discussed in more detail in 4.1.3).

FIGURE 5: Poverty in Household Income vs Welfare

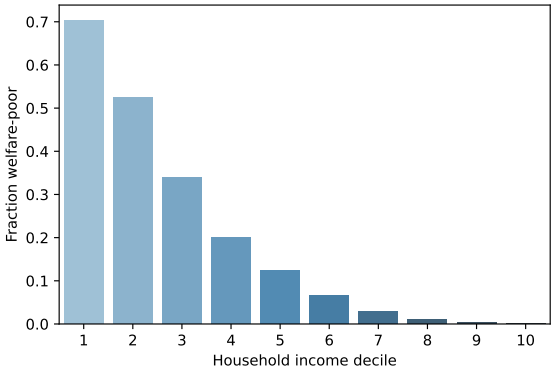

Notes: This figure shows the fraction of individuals who are welfare-poor according to the MMWI (i.e. in the lowest 20% of the distribution) for different deciles of the household income distribution.

Figure 5 shows the fraction of welfare-poor individuals in each of the deciles of the household income distribution. Under an income-based notion of poverty, all poor individuals would by definition be in the bottom two deciles. The figure shows that welfare-poor individuals are spread more broadly over the income distribution and can be found even in households above median income. The figure also shows that income certainly remains informative about an individuals' position in the welfare distribution. For example, in the bottom two income deciles, between 50-70% are also welfare-poor, whereas almost nobody is welfare-poor at the very top of the income distribution.

Table 8 summarises the relationship between poverty in household income and welfare-poverty by showing the extent to which both concepts agree in a 2x2 table. This table is computed on the level of individuals. This means that for example 72% of individuals are non-poor both according to household income and the MMWI. 12% are poor according to both welfare measures. Most importantly, there is a fraction of individuals where the measures give conflicting answers. 8% of the population are welfare-poor but not household-income-poor (*false negatives*) and 8% are household-income-poor but not welfare-poor (*false positives*). The numbers of false positives and false negatives are equal by construction as the quantity of poor individuals is kept constant across welfare measures (20%).

TABLE 8: Poverty in Household Income vs Welfare

	Not welfare-poor	Welfare-poor
Not HH-income-poor	0.72 (*True Negatives*)	0.08 (*False Negatives*)
HH-income poor	0.08 (*False Positives*)	0.12 (*True Positives*)

Notes: The table shows the classification of poor and non-poor individuals according to household income and welfare (i.e. the MMWI).

A useful summary statistic is the ability of household income to identify welfare-poor individuals in the population (the 'sensitivity' of household income in statistical terms).[20] Using household income as a welfare measure and defining poverty based on it allows to identify 61% of welfare-poor individuals ($= 0.12/0.2$). Conversely, this means that an analysis based on household income misses 39% of welfare-poor. This is an important result as it suggests that focusing on households below the income poverty threshold may lead to neglecting a substantial fraction of those are welfare-poor.

A related statistic is the ability of household income to identify the right individuals as non-poor (the 'specificity'). This statistic is relatively high, meaning that if somebody is not poor according to household income, then the likelihood that this person is not welfare-poor is 90%. This high value reflects that the unconditional probability of being poor is low according to each measure (20% by definition). As a result, if somebody is not household-income-poor, they come from the larger part of the income distribution and it is also relatively likely that this person is not welfare-poor.

4.1.2 What explains the difference between income and welfare poverty?

Table 8 has shown that household income does not identify all welfare-poor individuals in the population. An important question is what is driving the differences between household income and welfare. As highlighted earlier, there are three different components of individual welfare that are included in the MMWI but ignored by household income. These are (1) the unequal sharing of consumption within families (2) time use of each family member and (3) preferences. The idea behind the following exercises is to account for these additional aspects one-by-one in order to understand the role of these factors.

A first reason for why income does not identify all welfare-poor individuals is that couples can share consumption unequally. As a result, if we observe a couple with an income above the

[20] Sensitivity is defined as $\frac{TP}{TP+FN}$, where TP is the number of true positives and FN the number of false negatives. Specificity is $\frac{TN}{TN+FP}$, where TN is the number of true negatives and FP the number of false positives.

TABLE 9: Sensitivity and Specificity for Different Welfare Measures

Chosen Welfare Measure	Sensitivity	Specificity
(1) Income	0.61	0.90
(2) Consumption	0.64	0.91
(3) Empirical sharing rule	0.81	0.95

Notes: The table shows the sensitivity and specificity of using poverty according to different welfare measures if the 'true' welfare measure is the MMWI. Also see footnote 20 for the definition of these statistics.

poverty threshold, it can well be that only one partner is below the poverty line for consumption while the other partner is above. An easy way to construct a welfare index which improves on income and takes unequal consumption sharing into account is to compute *total individual consumption* (C_i) as the sum of personal consumption and money expenditure for the public good:

$$C_i = c_i + q_i$$

If the fact that income poverty does not identify all welfare-poor individuals was primarily driven by unequal consumption sharing, we would expect that defining poverty based on individual consumption (i.e. classifying the lowest 20% of the distribution of C_i) should do a much better job than income at identifying the welfare-poor. Table 9 shows that sensitivity does improve, as consumption-poverty can identify 64% instead of 61% of welfare-poor individuals. However, sensitivity is still far from 100%, which suggests that unequal sharing alone cannot explain the difference between income and welfare poverty.

The second mechanism that can be assessed is the role of including time use. Another reason why income might not identify all welfare-poor individuals is that it does not take time use (i.e. leisure and hours spent on home production) into account, which also affects welfare and is included in the MMWI. A straightforward way to account for time use when analysing welfare is to compute the so-called *empirical sharing rule (ESR)* (Lise and Yamada (2019)). This measure adds the market value of leisure and home production to total consumption of the individual. The market value is defined by using the wage rate of the individual, which measures the opportunity cost of time:

$$\text{ESR}_i = c_i + q_i + w_i l_i + w_i d_i + w_j d_j$$

Here, the index j refers to the partner of the individual. As home hours are public within the

household, the home hours of the partner also matter for the welfare of the individual. The empirical sharing rule accounts for unequal consumption sharing and time use. However, it still ignores preference heterogeneity as it is simply the sum of all goods, not weighted by how much an individual values the goods. For example, public good expenditure (q) contributes to the empirical sharing rule of both partners in the same way, even when their preferences differ. Still, using the empirical sharing rule leads to a sensitivity of 81%. This is a substantial improvement over welfare indices based on income or consumption and suggests that the empirical sharing rule is a better proxy for welfare than these welfare indices.

Finally, the role of preference heterogeneity can be assessed by looking at the fraction of welfare-poor individuals who are not identified by the empirical sharing rule. As indicated in the previous discussion, there are three aspects that income does not take into account which are included in the MMWI: (1) unequal sharing within the household (2) time use and (3) preference heterogeneity. As incorporating the first two of the factors into the analysis does not fully explain the difference (i.e. sensitivity stays below 100%), this implies that preference heterogeneity is important, which is the remaining factor that can lead to a difference between income and welfare poverty. In other words, accounting only for unequal sharing and time use leaves 19% of welfare-poor individuals unidentified.[21]

From an empirical perspective, note that both individual consumption and the empirical sharing rule are more straightforward to measure than the MMWI, as they don't require information on preferences. Measuring individual consumption requires a breakdown of consumption into public and private consumption, which is available for example in the datasets used in Lise and Yamada (2019) or Cherchye, De Rock, and Vermeulen (2012), and computing the empirical sharing rule requires also having information on time use. The results from the model suggest that while these more easily observable welfare measures can represent an improvement over household income, they also do not identify all welfare-poor individuals, highlighting the importance of also accounting for preferences.

[21] As an alternative assessment of the role of preferences, I also computed the MMWI under the assumption that everybody has identical preferences which are given by the sample mean of preferences. This hypothetical MMWI under homogeneous preferences is another way of taking unequal sharing and time use, but not preference heterogeneity into account. Defining poverty according to this measure identifies only 54% of welfare-poor individuals.

TABLE 10: Mean preferences of welfare and income-poor individuals

	General Population	Income-poor	Welfare-poor	Welfare-poor, not income-poor
α^c	0.19	0.12	0.21	0.28
α^l	0.45	0.52	0.44	0.39
α^D	0.36	0.35	0.35	0.33
α^c (Partner)	0.19	0.12	0.22	0.31
α^l (Partner)	0.45	0.53	0.48	0.44
α^D (Partner)	0.36	0.35	0.30	0.26
Economies of Scale	0.39	0.38	0.23	0.15

Notes: This table shows the mean characteristics in terms of wages and preferences of the individual and their partner for different groups. In particular, the last column shows the characteristics of those who are poor according to the MMWI but not according to household income.

To get a better picture of the reasons why household income does not fully capture welfare, it is also useful to have a closer look at the characteristics of the people who are welfare-poor but not poor in terms of income. These are shown in Table 10, which reports preferences for each good (α^c, α^l, α^D) and a measure of the economies of scale of the household.[22] The first column shows the mean of these variables for the population and the other columns allow to compare different groups of poor individuals. Compared to the population, income-poor individuals on average have a higher leisure preference, reflecting that this leads to lower working hours and thereby income. This is different for welfare-poor individuals, whose leisure preference is only slightly higher than for the population. Welfare-poor individuals also have a higher consumption and lower home good preference. Most interestingly, the table shows preferences for those who are welfare-poor but not income-poor. The striking feature of this column is that this group has the lowest leisure preference and the consumption preference, both for the individual and their partner, is highest. As a result, as the last row shows, this group has lower economies of scale as a result of their preferences. As there are more couples in which both partners value private consumption strongly, these couples can share fewer goods. These couples have an income that puts them above the poverty line, but as they are limited in their sharing, they are still welfare-poor. The table demonstrates that one way through which preference heterogeneity affects welfare-poverty is through economies of scale.

[22] Economies of scale are computed as $\frac{MMWI_i + MMWI_j - w_i - w_j}{w_i + w_j}$. Recall that the MMWI can be expressed in terms of an hourly wage.

TABLE 11: Number of poor individuals in the household

	Nobody is poor	1 person is poor	2 people are poor
(a) General Population			
Income	80.0	0.0	20.0
MMWI	74.1	12.1	13.9
(b) Non-income poor			
MMWI	85.8	9.0	5.2

Notes: The table reports the fraction of couples in which population in which 0/1/2 household members are poor. Section (a) shows these statistics for the population and section (b) for the group of households that are not income-poor. All values are reported in %.

4.1.3 Intra-household inequality

This section zooms in on intra-household inequality. As welfare-poverty is defined on the level of individuals, it is not necessarily the case that both household members are poor, but there can be cases when only one individual is poor, while their partner is non-poor. In section (a) of Table 11 , the table reports the breakdown of couples in the population by the number of poor individuals. For comparison, the table first shows the case of income-poverty, where by definition there are 20% of households in which both people are poor. The table demonstrates that the picture is fairly different for welfare-poverty. According to the MMWI, there is a substantial fraction (12%) of households in which only one person is poor, which can only be captured by an individual-level welfare analysis. In addition, there are 14% of households in which both members are poor. Note that the number of households that are affected by poverty grows relative to household income: according to the MMWI, there are 26% (instead of 20%) of families in which at least one household member is poor.

Section (b) of Table 11 shows a similar breakdown, but only for the case of households which are not income-poor. This sheds further light on the circumstances under which individuals can be welfare-poor, but not income-poor. Among the non-income-poor, there are 9% of households in which only one individual is welfare-poor, but their partner is not. In these cases, welfare is unevenly distributed within the household such that only member is below the poverty line according to the MMWI. In addition, there are 5% of households in which *both* household members are welfare-poor, despite them being above the poverty line for income. In these cases, welfare-poverty is not driven by intra-household sharing, but by the fact that given their household income, both household members obtain relatively low levels of welfare, for example because of

TABLE 12: Fraction of couples in which welfare measures agree

	MMWI
Consumption	0.48
Empirical Sharing Rule	0.59

Notes: This table shows the fraction of couples in which different welfare measures agree on who in the couple is better off. For example, the entry in the first row contains the fraction of couple for which individual consumption and the MMWI lead to the same welfare conclusion.

low economies of scale due to a high preference for private consumption (also see Table 10).

Table 12 looks at intra-houshold inequality from a different angle and asks what information is required to determine whether one household member is better or worse off than the other. This goes back to the initial example of Alex and Bertie. Suppose we observe that Bertie consumes more than Alex, can we conclude that Bertie is better off? If we also include their time use, but not preferences, can we conclude who is better off? The table reports the fraction of cases where the conclusion about who is better off (partner A or partner B) according to different welfare measures lead to the same conclusion that would be reached using the MMWI. Individual consumption leads to the same conclusion only in 48% of cases. This a slightly worse predictive performance than a random guess about which partner is better off, which would be correct in 50% of cases, meaning that just comparing consumption across partners is not very informative about the welfare ranking.

When comparing the welfare of partners with the empirical sharing rule, the common parts (i.e. public consumption and home hours) drop out, so that the empirical sharing rule compares the sum of consumption and leisure:

$$\text{ESR}_A > \text{ESR}_B \iff c_A + w_A l_A > c_B + w_B l_B$$

With the empirical sharing rule, the predictive power improves and the welfare conclusion is the same as according to the MMWI in 59% of cases. Thus, comparisons according to the empirical sharing rule reflect the underlying welfare ranking better than consumption. However, there is still a substantial gap which implies that it is essential to take preferences into account when making welfare comparisons between partners.

TABLE 13: Experiment: Eliminating intra-HH inequality and economies of scale

	Poverty rate (MMWI)	Inequality (MMWI)
Baseline	20.0	0.23
Equal consumption sharing	24.4	0.28
Equal sharing (MMWI)	19.1	0.19
No economies of scale	39.4	0.17

Notes: The table reports poverty rates for the baseline simulation (row 1) and several thought experiments (remaining rows), where intra-household inequality and economies of scales are eliminated. Poverty rates are reported in % and inequality as the value of $\text{Var}(\text{Log}(MMWI_i))$.

4.1.4 Eliminating intra-household inequality and economies of scale

The previous sections have highlighted the role of intra-household inequality and economies of scale in determining welfare-poverty. As a result, an interesting thought experiment is what would happen if each of these factors were eliminated. This sheds some further light on the role of intra-household inequality and economies of scale. In all of these experiments, the threshold for welfare-poverty is kept at its level in the baseline simulation, where it is computed such that 20% of the population are poor.

First, I focus on the role of intra-household inequality in private consumption. Suppose private consumption in each household was equalised by setting each individual's level to the household mean, while keeping all other choices of household and the stock of couples constant. Rather than decreasing welfare-poverty, Table 13 shows that this experiment *increases* the fraction of welfare-poor individuals. The reason for this effect is that the model allows for two sources of consumption differences within households. The first is unequal sharing (via the Pareto weight λ). If this was the primary mechanism, eliminating intra-household inequality would reduce welfare-poverty, as it would reallocate from partners with higher bargaining power within the household to those with lower bargaining power. However, the estimation suggests that a large part of consumption differences is driven by preference heterogeneity. In this case, eliminating consumption differences does not reduce welfare-poverty, as it does not necessarily reallocate consumption from the better-off to the worse-off partner. This experiment demonstrates that preference heterogeneity is important for the welfare consequences of consumption inequality.

Second, I study the role of eliminating intra-household inequality in welfare rather than consumption. I keep the stock of couples constant and compute the allocation that would lead to equal welfare for both partners for each couple. This is done by adjusting the Pareto weights

TABLE 14: Eliminating intra-HH inequality: Flows in and out of poverty

	Fraction of all households
Poor to non-poor	3.5
Non-poor to poor	2.7
Net effect on poverty rate	-0.9

Notes: The table reports the flows in and out of welfare-poverty when welfare is equalised within couples.

(λ_{it}) such that the MMWI of partner A is equal to the MMWI of partner B.[23] Table 13 shows that this results in a reduction of welfare-poverty. The poverty rate declines from 20% to 19.1%. The reason for the modest size of this effect is that there are two opposing effects on the poverty rate. The first effect ("poor to nonpoor" in Table 14) is that there are 3.6% of individuals in the population who are welfare-poor in the baseline simulation, but who become non-poor when welfare within couples is equalised. These individuals were worse off than their partner before, so that the reallocation pushes them above the poverty threshold. This is a fairly substantial effect, as it is effectively 18% of all poor individuals (= 3.6%/20%). However, there is a second effect ("nonpoor to poor" in the table). In these cases, reallocation towards the worse off partner does not lift them above the poverty line. Instead, *both* partners are poor after the reallocation, whereas before at least one person was above the welfare poverty line. Table 14 shows that this applies to 2.7% of all individuals in the population, resulting in the net effect of -0.9% reported in Table 13.

Taken together, while there is a substantial fraction of households in which only one partner is welfare-poor (Table 11), eliminating intra-household inequality only has a modest effect on reducing poverty rates as many of these families have too few resources for both partners to be above the welfare poverty line. To assess the effect of intra-household inequality on the wider distribution of welfare, Table 13 also shows how total inequality (measured by $\text{Var}(\text{Log}(MMWI_i))$ is affected. Eliminating intra-household inequality reduces total inequality by 16%, which is an important fraction of total inequality.

Finally, I eliminate economies of scale (EOS) by rescaling each individuals' MMWI by the measure of economies of scale (see footnote 22 for the definition). This experiment has a drastic impact on the poverty rate, as it almost doubles to 39%. This experiment highlights the role of EOS for the assessment of poverty. Interestingly, eliminating EOS substantially *reduces* inequality (column 2), as the variance decreases by 25%. This reflects that EOS are unevenly distributed in the

[23]Note that this does not mean that the Pareto weight is 0.5 for each couple. $\lambda = 0.5$ only corresponds to welfare equality when partners have equal preferences and equal wages. Figure 10 shows the distribution of Pareto weights in this case.

population. In some couples, both partners value the home good strongly and can share a lot of resources, while other couples value leisure and private consumption more strongly and can share less. Economies of scale therefore have the double role of reducing poverty while increasing inequality.

4.2 Policy Simulation: Minimum Wages and Poverty

4.2.1 Policy Background

In 2021, the minimum wage in the UK is £8.91 for workers aged 23 and over.[24] According to Dube (2019), around 7% of the workforce earns the minimum wage. A well-known fact is that many minimum wage earners live in households which are above the poverty threshold, for example in couples where only the secondary earner receives the minimum wage (Burkhauser (2015)). Panel (a) of Figure 6 illustrates this point in the context of the UK. Cribb, Joyce, and Xu (2019) compute the fractions of minimum employees which live in households in different deciles of the income distribution. The striking feature of this graph is that minimum wage earners are found in all deciles of the income distribution and the peak is the 6th decile, with around 14% of minimum wage earners being in this category. Panel (b) of Figure 6 from Cribb, Joyce, and Norris Keiller (2017) looks at the fraction of the mechanical gains in household income from concrete policy proposals which accrues to households in different income deciles. This graph mirrors panel (a). Since many minimum wage earners live in households well above the poverty threshold, a large fraction of the income gain (assuming no behavioural changes) goes to non-poor households. These figures underscore the point by Burkhauser (2015) that many minimum wage earners do not live in poor families.

4.2.2 Experiment

From an intra-household perspective, there is the additional intricacy that just because a minimum wage earner does not live in an income-poor household, they can still be poor in terms of their individual well-being. This suggests that increases in the minimum wage might reduce individual-level welfare poverty even among those minimum wage earners who live in a non-poor household. Figure 7 shows poverty rates according to the MMWI among minimum wage earners, depending on the wage of the partner of the individual. The figure shows that the fraction of poor individuals is very high when both partners are minimum wage earners (group 1). However, poor individuals can be found in all groups, also in those where the partner has high wages (groups 3 and 4). This captures the same intuition as Figure 5. The main difference is the breakdown by the wage of the partner, rather than household income. This will be useful in the

[24] There are lower levels are younger workers.

FIGURE 6: Minimum Wage Earners along the Income Distribution in the UK

(a) Fraction of MW earner by income decile

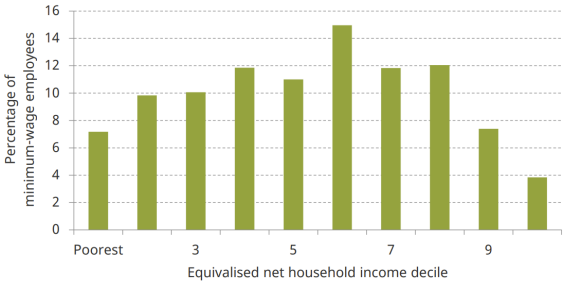

(b) Fraction of HH income increase by decile

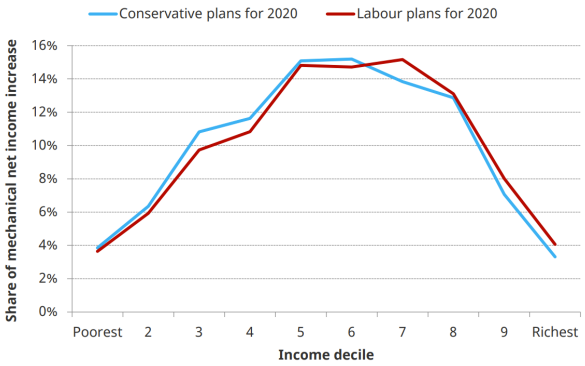

Notes: Panel (a) (from Cribb, Joyce, and Xu (2019)) shows the fraction of employees who receive the minimum wage along the income distribution based on data from the Family Resources Survey and the Labour Force Survey. Panel (b) (from Cribb, Joyce, and Norris Keiller (2017)) shows the fraction of the mechanical income gains from the Conservative and Labour proposal which accrue to different deciles of the income distribution.

policy experiment, as the wage distribution is exogenous and can be compared more easily across policy regimes than the endogenous income distribution.

I perform an experiment with a minimum wage increase, by raising the lowest wage in the economy exogenously by 20%. Note that I abstract from supply-side employment effects. The goal of the analysis is to study the impact of the minimum wage increase on minimum wage earners and their partners, and how the benefits from the increase are shared within families.[25] Table 15 shows the impact of the reform on different outcomes, focusing on minimum wage earners only and breaking the analysis down by the different wage types of their partners. Note that I compare welfare between steady states, which can be interpreted as the long-run impact of the minimum wage increase once the economy has transitioned to the new marriage market equilibrium.

Panel (b) of Figure 7 shows the impact of the minimum wage increase on poverty rates, depending on the wage group of the partner. The main result from this analysis is that welfare poverty decreases in all groups, also when the partner of the minimum wage earner has a relatively high wage. For example, in couples where both individuals are minimum wage earners (group 1), the reform reduces welfare poverty by 6 pp. In the other groups, welfare poverty among minimum wage earners decreases by 4-7 pp, even when the partner has a substantially higher wage (such as on average £22.3 in group 3). These results imply that the minimum wage does not only decrease poverty in couples where both partners have a low wage, but also in couples where the minimum

[25] In a review of the available evidence, Dube (2019) concludes that such employment effects are likely to be small, at least for moderate minimum wage increases. It would be possible to include supply-side effects in this model, for example by having a type-specfic frictionless labour market with market clearing, but this would make the model more complicated.

FIGURE 7: Minimum Wage (MW) Analysis - Main Results

(a) Welfare-poor MW earners by partners' wage

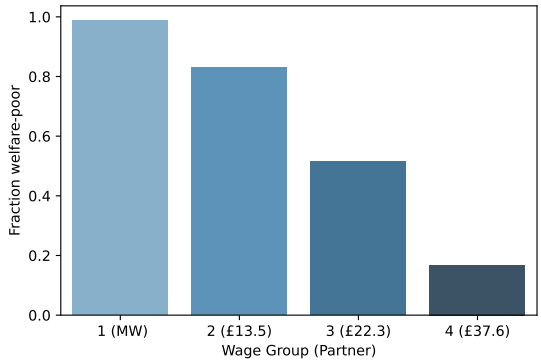

(b) Reduction of welfare poverty by partners' wage

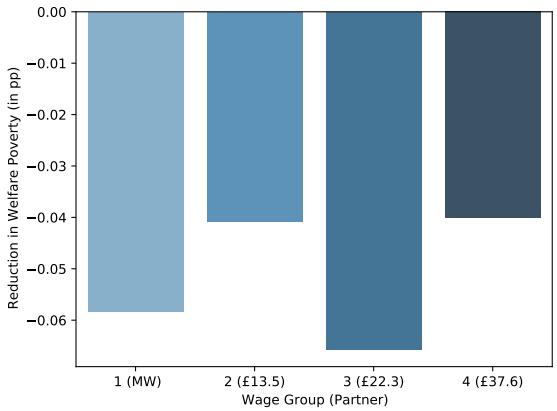

Notes: Panel (a) shows the fraction of individuals who are welfare-poor (i.e. in the lowest 20% of the distribution of the MMWI) for minimum wage earners with different groups of partner earnings. In group 1, the partner is also a minimum wage earner. For the other groups, the average earnings of the partner are given by the number in brackets. Panel (b) shows the reduction in poverty rates that results from the minimum wage increase. The poverty thresholds in the experiment are kept constant at their values in the baseline simulation (panel a).

wage earner has a higher earning partner, but resources are shared unequally within the couple.

Table 15 shows the reduction in poverty rates along with the impact of the reform on income, average MMWI and the Pareto weight within couples. Naturally, the impact on household income is highest for those couples where both partners are minimum wage earners and lowest when the partner has the highest wage type. The change in average welfare (measured by the MMWI) is fairly similar to the change in household income, and also similar for both partners. In addition, there is also an increase in the Pareto weight of the minimum wage earner. The reason for this that the policy increases the relative wage of minimum wage earners and thereby their intra-household decision power.[26] The only exception is the case of couples where both partners are minimum wage earners. As relative wages are not affected in this case, the Pareto weights remains unchanged.

[26]Formally, in the model, Pareto weights are linked to the relative value of singlehood and these values change in favour of minimum wage earners after the reform.

TABLE 15: Impact of experiment on MW earners (by partners' wage type)

Wage Type (Partner)	1 (MW)	2 (£13.5)	3 (£22.3)	4 (£37.6)
Household Income (in %)	0.17	0.09	0.06	0.03
MMWI (in %)	0.21	0.08	0.05	0.03
MMWI (Partner, in %)	0.21	0.11	0.04	-0.00
Welfare Poverty (in pp)	-0.06	-0.02	-0.06	-0.04
Pareto Weight (λ)	0.00	0.02	0.03	0.02

Notes: The table reports the impact of the minimum wage increase on household income and the MMWI (scaled in % changes), individual-level poverty (scaled in percentage points) and the Pareto weight in the couple (scaled between 0 and 1).

4.2.3 Decomposition of mechanisms

To better understand the mechanisms behind the poverty reduction due to the minimum wage increase, it is useful to consider a decomposition. This exercise breaks down the total reduction of poverty for each group (which is shown in panel (b) of Figure 7) into three distinct channels and allows to assess their relative importance.

The first of these steps is the *direct effect*. This corresponds to holding both time use choices and the distribution of couples constant. The minimum wage increase leads to an increase in the household budget and the couple can spend more on each partners' personal consumption and the public good. To allocate the additional budget across partners and the public good, I compute the fraction of total expenditure of each of these goods under the status quo and distribute the additional resources according to these shares.[27]

The second step is the *reoptimisation* effect. The minimum wage increases changes the budget constraint of households and couples reoptimise the time spent on leisure, work and home production. This step is important as changes in relative wages are likely to trigger changes in relative leisure, which affects individual-level welfare. This effect is computed by keeping the distribution of couples as well as Pareto weights constant and solving for the new optimal choices. This would be the impact of the reform in a 'unitary' version of the model in which Pareto weights are fixed.

Finally, the third step is the *bargaining effect*, which captures that the minimum increase could increase the decision power of minimum wage earners within the couple, as the relative wage

[27]For example, in a couple in which 30% is spent on partner A, 40% on partner B and 30% on the public good, the extra budget is spent accordingly.

shifts in their favour. In order to compute this effect, I compare the previous step with the new steady-state. Note that there can also be an effect through endogenous marriage and divorce, since the stock of married couples may look different in the new steady state.[28]

FIGURE 8: Poverty Reduction among MW earners - Mechanisms

(a) Decomposition

(b) Role of Bargaining Effect

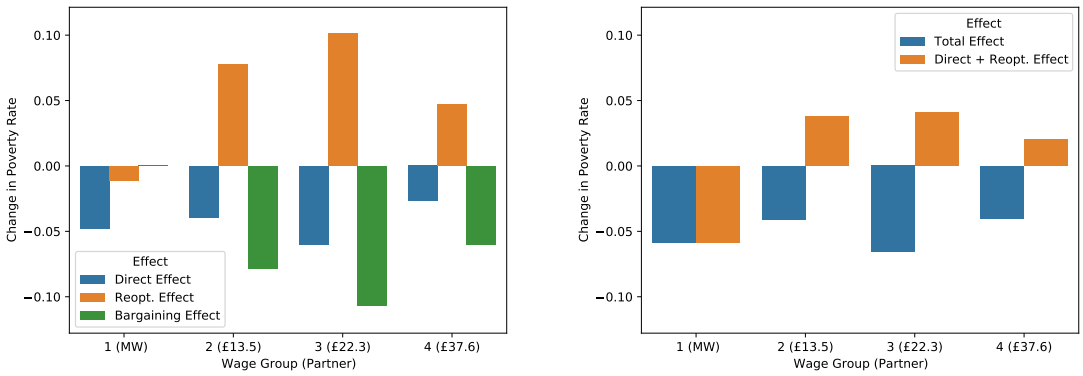

Notes: Panels (a) shows the decomposition of the total welfare poverty reduction into the direct, reoptimisation and bargaining effect. See main text for the definition of each of the effects. The sum of the three effects is equal to the total effect. Panels (b) focuses on the role of the bargaining effect (i.e. endogenous Pareto weights). This figure compares the total effect and the sum of the direct and reoptimisation effect (which ignores the bargaining effect).

The results from the decomposition are shown in Figure 8. First focus on panel (a), which shows each of the three effects. The sum of these effects add up to the total effect. As the direct effect increases the budget of the household, it reduces poverty across all of the groups. Note that the direct effect affects partners differentially, as the additional resources are not distributed equally across partners. The reoptimisation effect *increases* poverty in all couples where the partners earns more than the minimum wage. The mechanism behind this is that as the relative wage of the minimum wage earner increases, the opportunity cost of their leisure rises. As a result, the household increases their work hours and reduces their leisure, and this reduction in leisure raises poverty on the individual level. Quantitatively, the reoptimisation effect tends to overturn the direct effect in most cases. Finally, the bargaining effect allows to quantify the welfare impact of changes in the Pareto weight. This reduces poverty again, as it increases consumption, leisure and decision power over the public good. The bargaining effect tends to offset the reoptimisation effect, so that the total effect is close to the direct effect. Overall, all three effects play a role. In a similar spirit, Figure 11 shows the direct and the reoptimisation effect for different deciles of the income distribution, which further illustrates how these effects operate all across the distribution.[29]

[28]The distinction between the reoptimisation and the bargaining effect is related to Knowles (2012), who studies aggregate labour supply and contrasts the implications of a unitary and collective model.

[29]Note that it is not possible to produce a similar figure for the collective effect, as the stock of couples changes

The interpretation of these results is that the poverty reduction shown in Figure 7 is primarily driven by the direct effect. For a given level of household income, the welfare of each partner can be different, so that only one person can be welfare-poor while the other is not. As the direct effect gives more resources to the household and increases the budget, more money can be spent on the consumption of each partner as well as the public good. This increases the welfare-level of both partners and the less well-off partner, who was welfare-poor before the reform, is pushed above the poverty threshold. The reoptimisation and bargaining effect are also quantitatively import, but cancel each other out, so that the total effect is close to the direct effect.

An important implication of these results is that the change in Pareto weights has significant welfare impacts. This suggests that it would be misleading to analyse the impact of minimum wage changes on 'unequal' couples (where one partner has a higher wage than the other) using a so-called unitary model of the household, where the Pareto weight is an exogenous parameter that is unaffected by policy changes. Panel (b) show the role of the bargaining effect by comparing the total effect and the effect that would be obtained when ignoring the bargaining effect. These figures show that the bargaining effect does not play much of a role of partners have similar wages (i.e. when both are minimum wage earners), but that the predictions between a unitary and collective version of the model are fairly different as soon as wages differ. In several of these cases, a unitary model would imply that a minimum wage increase would on average *harm* minimum wage earners and increase their poverty if they live with a higher-wage partner.

between the pre- and post-reform simulations. As a result, the figures in the main text have the *wage* of the partner on the x-axis.

5 Conclusion

In this paper, I have studied the link between household income and individual welfare and the implications for measuring poverty at the level of individuals. Individual welfare is a broader concept that takes unequal sharing, time use and preferences into account. In particular, preference heterogeneity is an important consideration for two reasons. First, preference heterogeneity matters for interpreting observed differences in consumption and leisure between partners. These could reflect either inequality (through differences in the Pareto weight) or preference heterogeneity. For example, observing that one person consumes more than their partner does not necessarily imply that this person is better off, but it could also reflect preferences. Second, accounting for preference heterogeneity allows for a more nuanced view on public goods. Even though both partners have access to the public good to the same degree, they can have different valuations for it and may not benefit to the same extent. These different valuations need to taken into account to measure the living standard of the individual. To study these issues and their role for the welfare analysis, I build a structural model of time use and the marriage market. The model allows to pin down the joint distribution of unobserved preferences of partners through a combination of data and theory, by using data on the variability of choices among singles along with a model of the marriage market.

I find that poverty in household income is an imperfect predictor of individual-level welfare poverty. Welfare poverty is based on considering an individuals' total utility from consumption, leisure and the home good. Income poverty only identifies 61% of welfare-poor individuals, depending on the measure of individual welfare which is used. I investigate the role of unequal sharing, time use and preferences separately and find that preference heterogeneity plays an important role. Accounting for the first two channels without taking preferences into consideration only allows to identify 81% of welfare poor individuals, which implies that accounting for preference heterogeneity is important to detect the remaining 19% of the welfare-poor.

I highlight the policy relevance of using individual welfare measures by studying a hypothetical increase in the minimum wage. It has been documented that many minimum wage earners live in non-income-poor families. As a result, according to a conventional view of poverty based on household income, minimum wage increases might seem not well targeted to reduce poverty. In my model, minimum wage earners can still be poor even if they live in a non-income-poor family. As a result, minimum wage increases can decrease individual poverty even in families who are not income poor.

These results demonstrate that an individual welfare analysis can lead to different conclusions than a conventional analysis based on household income. Perhaps most importantly, this calls for more data collection on each of the elements in the link between income and welfare in order to develop robust empirical measures of individual welfare. My results suggest that data on individual consumption, time use and preferences is required in order to identify welfare-poor individuals. There are measurement issues related to each of these points that should be

addressed in future research. For many countries, including the UK, there are no datasets that simultaneously measure consumption and time use.[30] In addition, there are several methodological questions that could be addressed, such as reducing measurement error in retrospective questions about consumption or time use or better measuring the breakdown of consumption into personal and public goods. Finally, no direct data on the type of preferences needed in collective household models (the relative valuation of consumption, leisure and home goods) is available. An important question is whether and how these preferences can be measured directly, as it is less clear if they can be reliably quantified through hypothetical survey questions. The results from this paper suggest that advances in this direction will lead to a better understanding of individual-level welfare.

The are several other important directions for future research which would shed further light on the individual-level distribution of welfare. One simplifying assumption in this paper was that there is a single household-level public good which individuals have different preferences over. In reality, there are several public goods within households and it is likely that partners often disagree about how to allocate expenditure across public goods (such as money spent on children, cars, furniture or gardening). Studying preference heterogeneity on this more fine-grained level would allow to make a better distinction to what extent partners benefit equally from public expenditure or whether partners' valuation of these goods diverges substantially. In addition, there are dynamic considerations, such as intra-household differences in patience or risk aversion, that this paper has abstracted from. This would allow to assess to what extent e.g. assets, human capital or portfolio choices are valued equally by both partners, and could be incorporated into the welfare measures.

[30] The Dutch LISS panel or the Japanese Panel Survey of Consumers are examples of such datasets.

References

Atkinson, Anthony B and François Bourguignon. 2014. *Handbook of income distribution*, vol. 2. Elsevier.

Bargain, Olivier, Olivier Donni, and Prudence Kwenda. 2014. "Intrahousehold distribution and poverty: Evidence from Cote d'Ivoire." *Journal of Development Economics* 107:262–276.

Blundell, Richard, Pierre-André Chiappori, and Costas Meghir. 2005. "Collective labor supply with children." *Journal of political Economy* 113 (6):1277–1306.

Bronson, Mary Ann and Maurizio Mazzocco. 2019. "Taxation and Household Decisions: an Intertemporal Analysis." *Unpublished manuscript* .

Brown, Caitlin, Rossella Calvi, and Jacob Penglase. 2021. "Sharing the pie: An analysis of undernutrition and individual consumption in Bangladesh." .

Browning, Martin, Laurens Cherchye, Thomas Demuynck, Bram De Rock, and Frederic Vermeulen. 2021. "Stable marriage, household consumption and unobserved match quality." *Available at SSRN 3918838* .

Browning, Martin, Pierre-Andre Chiappori, and Arthur Lewbel. 2013. "Estimating consumption economies of scale, adult equivalence scales, and household bargaining power." *Review of Economic Studies* 80 (4):1267–1303.

Burkhauser, Richard. 2015. "The minimum wage versus the earned income tax credit for reducing poverty." *IZA World of Labor* .

Calvi, Rossella. 2020. "Why are older women missing in India? The age profile of bargaining power and poverty." *Journal of Political Economy* 128 (7):2453–2501.

Cherchye, Laurens, Sam Cosaert, Bram De Rock, Pieter Jan Kerstens, and Frederic Vermeulen. 2018. "Individual welfare analysis for collective households." *Journal of Public Economics* 166:98–114.

Cherchye, Laurens, Bram De Rock, Arthur Lewbel, and Frederic Vermeulen. 2015. "Sharing rule identification for general collective consumption models." *Econometrica* 83 (5):2001–2041.

Cherchye, Laurens, Bram De Rock, Khushboo Surana, and Frederic Vermeulen. 2020. "Marital matching, economies of scale, and intrahousehold allocations." *Review of Economics and Statistics* 102 (4):823–837.

Cherchye, Laurens, Bram De Rock, and Frederic Vermeulen. 2012. "Married with children: A collective labor supply model with detailed time use and intrahousehold expenditure information." *American Economic Review* 102 (7):3377–3405.

Cherchye, Laurens, Thomas Demuynck, Bram De Rock, and Frederic Vermeulen. 2017. "Household consumption when the marriage is stable." *American Economic Review* 107 (6):1507–34.

Chiappori, Pierre-André, Monica Costa Dias, and Costas Meghir. 2018. "The Marriage Market, Labor Supply, and Education Choice." *Journal of Political Economy* 126 (S1):S26–S72.

Chiappori, Pierre-Andre and Maurizio Mazzocco. 2017. "Static and intertemporal household decisions." *Journal of Economic Literature* 55 (3):985–1045.

Chiappori, Pierre-André and Costas Meghir. 2014. "Intra-household welfare." Tech. rep., National Bureau of Economic Research.

———. 2015. "Intrahousehold inequality." In *Handbook of Income Distribution*, vol. 2. Elsevier, 1369–1418.

Ciscato, Edoardo. 2019. "The Changing Wage Distribution and the Decline of Marriage." *Unpublished manuscript* .

Ciscato, Edoardo, Alfred Galichon, and Marion Goussé. 2020. "Like attract like? A structural comparison of homogamy across same-sex and different-sex households." *Journal of Political Economy* 128 (2):740–781.

Cribb, J, R Joyce, and A Norris Keiller. 2017. "Minimum Wages in the next Parliament." In *IFS Election Briefing Note*, vol. 205.

Cribb, J, R Joyce, and Xiaowei Xu. 2019. "The future path of minimum wages." *IFS Briefing Note, BN260, Institute for Fiscal Studies, November, https://www. ifs. org. uk/election/2019/article/the-future-path-of-minimum-wages* .

Decoster, André MJ and Peter Haan. 2015. "Empirical welfare analysis with preference heterogeneity." *International Tax and Public Finance* 22 (2):224–251.

Doepke, Matthias and Michèle Tertilt. 2016. "Families in macroeconomics." In *Handbook of macroeconomics*, vol. 2. Elsevier, 1789–1891.

Dube, Arindrajit. 2019. "Impacts of minimum wages: review of the international evidence." *University of Massachusetts, Amherst. https://www. gov. uk/government/publications/review-of-theinternational-evidence-on-the-impacts-of-minimum-wages/review-of-the-international-evidenceon-the-impacts-of-minimum-wages-terms-of-reference* .

Dunbar, Geoffrey R, Arthur Lewbel, and Krishna Pendakur. 2013. "Children's resources in collective households: identification, estimation, and an application to child poverty in Malawi." *American Economic Review* 103 (1):438–71.

Fields, Gary S and Ravi Kanbur. 2007. "Minimum wages and poverty with income-sharing." *The Journal of Economic Inequality* 5 (2):135–147.

Fleurbaey, Marc. 2006. "Social welfare, priority to the worst-off and the dimensions of individual well-being." *Inequality and Economic Integration, London: Routledge* .

Fleurbaey, Marc and François Maniquet. 2011. *A theory of fairness and social welfare*, vol. 48. Cambridge University Press.

———. 2018. "Optimal income taxation theory and principles of fairness." *Journal of Economic Literature* 56 (3):1029–79.

Fleurbaey, Marc et al. 2008. *Fairness, responsibility, and welfare*. Oxford University Press.

Foerster, Hanno. 2020. "Untying the Knot: How Child Support and Alimony Affect Couples' Decisions and Welfare." Tech. rep., University of Bonn and University of Mannheim, Germany.

Gayle, George-Levi and Andrew Shephard. 2019. "Optimal taxation, marriage, home production, and family labor supply." *Econometrica* 87 (1):291–326.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L Violante. 2010. "The macroeconomic implications of rising wage inequality in the United States." *Journal of political economy* 118 (4):681–722.

Knowles, John A. 2012. "Why are married men working so much? An aggregate analysis of intra-household bargaining and labour supply." *Review of Economic Studies* 80 (3):1055–1085.

Krueger, Dirk and Fabrizio Perri. 2006. "Does income inequality lead to consumption inequality? Evidence and theory." *The Review of Economic Studies* 73 (1):163–193.

Lechene, Valérie, Krishna Pendakur, and Alex Wolf. 2019. "OLS estimation of the intra-household distribution of consumption." Tech. rep., IFS Working Papers.

Lise, Jeremy and Shannon Seitz. 2011. "Consumption inequality and intra-household allocations." *The Review of Economic Studies* 78 (1):328–355.

Lise, Jeremy and Ken Yamada. 2019. "Household sharing and commitment: Evidence from panel data on individual expenditures and time use." *The Review of Economic Studies* 86 (5):2184–2219.

Low, Hamish, Costas Meghir, Luigi Pistaferri, and Alessandra Voena. 2018. "Marriage, Labor Supply and the Social Safety Net." *Unpublished manuscript* .

Lundberg, Shelly J, Robert A Pollak, and Terence J Wales. 1997. "Do husbands and wives pool their resources? Evidence from the United Kingdom child benefit." *Journal of Human Resources* :463–480.

Mazzocco, Maurizio, Claudia Ruiz, and Shintaro Yamaguchi. 2013. "Labor supply, wealth dynamics, and marriage decisions." *Unpublished manuscript* .

Obermeier, Tim. 2019. "The Marriage Market, Inequaliy and the Progressivity of the Income Tax." *CRC TR 224 Discussion Paper Series* .

Penglase, Jacob. 2021. "Consumption Inequality among Children: Evidence from Child Fostering in Malawi." *The Economic Journal* 131 (634):1000–1025.

Reynoso, Ana. 2018. "The impact of divorce laws on the equilibrium in the marriage market." .

Shephard, Andrew. 2019. "Marriage market dynamics, gender, and the age gap." *Unpublished manuscript* .

Tommasi, Denni. 2019. "Control of resources, bargaining power and the demand of food: evidence from PROGRESA." *Journal of Economic Behavior & Organization* 161:265–286.

Voena, Alessandra. 2015. "Yours, Mine, and Ours: Do Divorce Laws Affect the Intertemporal Behavior of Married Couples?" *The American Economic Review* 105 (8):2295–2332.

Appendix

Appendix A: The Role of the Normative Framework

The issue of interpersonal utility comparisons with preference heterogeneity has also been addressed by a broader literature in welfare economics, which has highlighted the ethical judgments underlying the choice of different welfare criteria (Fleurbaey (2006), Decoster and Haan (2015)), although typically in the context of single-agent models.[31][32] The goal of this section is to demonstrate that the MMWI can also be viewed through the framework of Fleurbaey (2006) and to introduce an additional welfare criterion, the Rente Criterion, which captures different ethical judgments. I then discuss to what extent the main results from the welfare analysis depend on the choice of the normative framework.

Fleurbaey (2006) defines the "Rente Criterion" as an individuals' answer to the question:

> "What income would be enough for you, in replacement of your current situation, if you did no longer have to earn it?"

In the model, the Rente Criterion can be computed similarly to the MMWI by solving a hypothetical problem for each individual, where the individual is single, does not work and receives non-labour income z. The indirect utility V_i^R is defined as:

$$\begin{aligned} V_i^R(z) &= \max_{c,q,l,d} u_i(c,l,D) \\ c+q &= z \\ d+l &= 1 \\ D &= D(C,d) \end{aligned}$$

Note that the indirect utility function depends on i, as the preference type will affect the solution. The Rente Criterion R_i is defined implicitly as the non-labour income which yields utility level u_i:

$$V_i^R(R_i) = u_i$$

Furthermore, Fleurbaey (2006) defines the "Wage Rate Criterion":

> "What net wage rate would be enough for you, in replacement of your current situation, if you could adjust your amount of work as you wished?"

[31] Also see e.g. Fleurbaey and Maniquet (2018), Fleurbaey et al. (2008) or Fleurbaey and Maniquet (2011).

[32] These single-agent models can be interpreted as either individuals or 'unitary' household models, where households have a stable utility function and household welfare is the main outcome of interest.

The Wage Rate Criterion is identical to the MMWI.[33] From this perspective, the MMWI is part of a broader family of welfare measures which take preference heterogeneity into account in different ways.[34] In general, note that the MMWI and Rente Criterion would lead to identical conclusions in the absence of preference heterogeneity, and that it is in particular the presence of endogenous labour supply that leads to the difference between the two.

Fleurbaey (2006) argues that these two criteria capture intrinsically different notions of justice. These are best illustrated by looking at an example in the context of singles. Table 16 shows the choices and welfare criteria for the three hypothetical individuals from the main text, where the first is consumption-oriented, the second leisure-oriented and the third home-oriented.

From the perspective of the MMWI, these three singles are equally well off, as their full income is the same independently of the choices they make. The MMWI is given by £10 in each case and there is no inequality between them. By contrast, if we were comparing the singles by their household income, the consumption-oriented single would be best off, as he or she decides to work longer hours than the others to obtain the highest income. The Rente Criterion closely mirrors household income. Recall that the Rente Criterion is defined as the hypothetical *nonlabour* income that would be needed to get the same utility while not working. Not having to work gives the individual more time for leisure and home production. As a result, the Rente Criterion is lower than household income, since a lower amount of nonlabour income is needed to achieve the same utility as while having to work. The Rente Criterion is also highest for the consumption-oriented single.

This example illustrates the underlying ethical judgments involved in the choice between the MMWI and Rente Criterion. The Rente Criterion is based on the optimal choice between consumption, leisure and the home good, taking preferences into account. The consumption-oriented single is as well off as if they had a relatively high non-labour income, whereas the leisure and home-oriented singles are as well of as having a relatively lower non-labour income. The view behind the Rente Criterion is that those with the lowest levels of monetary resources are worst off, irrespective of whether this results from preferences (i.e. their choice to work fewer hours). As a result, one would conclude that the latter two singles are worse off than the consumption-oriented single. The MMWI aims to isolate the inequality that stems from constraints (the wage), rather than from choices. This means that individuals are viewed as being responsible for their choices if these result from preference heterogeneity. As a result, the MMWI concludes that all three individuals are equally well off, even though their income levels are different.

[33] Chiappori and Meghir (2015) define the MMWI in terms of the expenditure function, by computing the expenditure the individual would need as a single in order to achieve the same utility. Note that this notion of expenditure includes time use by rewriting the budget constraint in terms of full income (using the fact that the time budget is normalised to 1): $c + q + wl + wd = w$. As a result, the expenditure of a single is equal to their wage, so that the Wage Criterion is similar to the MMWI.

[34] Fleurbaey (2006) also defines a third criterion, the "Rente + Minimum Wage Criterion", where individuals are allowed to work at the minimum wage while receiving a certain non-labour income. This measure is an intermediate case between the Rente and the Wage Criterion.

TABLE 16: Welfare Measures - Examples (Singles)

Preferred Good	Consumption (c)	Leisure (l)	Home (D)
(a) Income/Consumption			
Income	2516.62	2009.55	2252.93
(b) Time Use			
Market hours	60.4	48.23	54.07
Leisure	31.99	44.24	31.29
Home Hours	19.61	19.53	26.64
(c) Welfare			
MMWI	10.0	10.0	10.0
Rente Criterion	1454.08	1072.26	1221.44
Utility	-3.6	-3.63	-3.72

Notes: The table shows examples for the welfare measurement for three different singles with different preferences. Time use is reported in weekly hours, the MMWI in £ per hour and income and the Rente Criterion in £ per month.

Figure 9 first shows the breakdown of welfare-poor individuals (according to the MMWI and the Rente Criterion) by household income decile. The rates are comparable between the two criteria. For example, in the lowest HH income decile, around 82% of individuals are poor according to the Rente Criterion and 70% are poor according to the MMWI

While these rates are fairly similar, this does not mean that the criteria classify *the same* individuals as poor. The MMWI identifies only 60% of those who are poor according to the Rente Criterion (and vice versa, as shown in Table 18). This number is remarkable low and indicates substantial differences between the two criteria, highlighting the role of the ethical judgments underlying the choice between the two. In order to better understand these differences, it is useful to consider how poverty according to each measure correlates with individual characteristics (i.e. preferences and wages). Table 17 shows these correlations. Poverty according to the Rente Criterion correlates strongly with the preference for leisure. The reason for this is that the Rente Criterion is based on the non-labour income an individual would need to achieve a certain living standard. Individuals with a high leisure preferences do not depend on income as strongly and tend to need a lower income to achieve a given living standard. The MMWI correlates less strongly with preferences and more strongly with labour market ability. Overall, the table highlights that the

ethical judgments underlying the choice between the MMWI and the Rente Criterion highlight who is considered as poor.

Table 18 also shows the breakdown of the role of unequal sharing, time use and consumption for the Rente Criterion. Consumption-poverty is much more predictive for poverty according to the Rente Criterion than for the MMWI (i.e. the MMWI in the main text). Consumption-poverty has a sensitivity of 80%, as opposed to 66% for income-poverty. The reason for this is that low levels of consumption correlate with a low consumption preference, and this is turn leads to lower values of the Rente Criterion (as individuals who value consumption little need relatively little non-labour income to achieve their utility level). Looking at the empirical sharing rule and also taking time use into account *decreases* sensitivity (to 58%), as the empirical sharing does not correlate with low consumption preferences as much as consumption. The conclusion from these results is that preference heterogeneity also plays an important role poverty according to the Rente Criterion, although it is now consumption-poverty which is the best proxy if information on preferences are not available.

FIGURE 9: Welfare vs Income Poverty

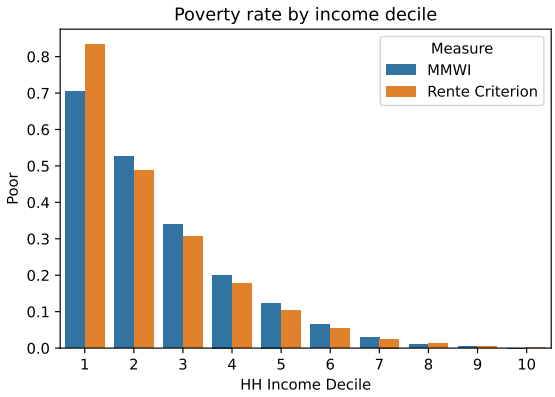

Notes: This figure shows the fraction of individuals who are poor according to the Rente Criterion and the MMWI (i.e. in the lowest 20% of the distribution) for different deciles of the household income distribution.

TABLE 17: Correlation between poverty, preferences and wages

Welfare Measure Characteristic	Rente Criterion	MMWI
Consumption pref. (α_i^C)	-0.34	0.07
Leisure pref. (α_i^L)	0.40	-0.02
Home Good pref. (α_i^D)	-0.04	-0.04
Ability (a_i)	-0.32	-0.46

Notes: The table shows how poverty according to the MMWI and the Rente Criterion correlate with individual characteristics (preferences and labour market ability).

TABLE 18: Sensitivity and Specificity for Different Welfare Measures

	Sensitivity		Specificity	
True Welfare Measure Chosen Welfare Measure	Rente Criterion	MMWI	Rente Criterion	MMWI
(1) Income	0.66	0.61	0.92	0.90
(2) Consumption	0.79	0.64	0.95	0.91
(3) Empirical sharing rule	0.58	0.81	0.89	0.95
(4) Rente Criterion	1.00	0.59	1.00	0.90
(5) MMWI	0.59	1.00	0.90	1.00

Notes: The table shows the sensitivity and specificity of using poverty according to different welfare measures if the 'true' welfare measure is either the Wage or the Rente Criterion. For example, the entry in the first row and first column describes the fraction of welfare-poor individuals according to the Rente Criterion who are identified by defining poverty based on income. Also see footnote 20 for the definition of these statistics.

Appendix B: Computational Details

Computing an equilibrium of the model requires solving a dynamic matching problem. This requires individuals to have rational expectations over future distributions of potential partners. As the expectation over who will be available in the future affects marriage market choices and thereby the future distributions, the equilibrium can only be solved for via fixed-point iteration. The general algorithm for solving the model proceeds as follows:

1. Pre-compute static choices conditional on all possible state variables. Pre-computing these choices is important from a computational perspective, as it avoids having to solve for optimal choices, which is nonlinear optimisation problem, during the rest of the model solution. Recall that the state vectors for couples and singles are given by:

$$\begin{aligned} \omega_i^S &= (g, a_i, p_i, b) \\ \omega_i^M &= (a_f, a_m, p_f, p_m, \lambda, b) \end{aligned}$$

   The optimal choices are computed for all combinations of the these state variables (using a grid for the Pareto weight λ and a standard algorithm such as BFGS):

$$\begin{aligned} (c, q, h, l, d) &= P^S(\omega^S) \\ (c_f, c_m, q, h_f, h_m, l_f, l_m, d_f, d_m) &= P^M(\omega^M) \end{aligned}$$

   During the rest of the model solution, optimal choices can then be obtained by interpolating this function linearly in the Pareto weight. Note that the Pareto weight is treated as a continuous state variables in the solution of the life-cycle problem and simulation.

2. Make a guess for the distributions of potential partners of each gender and age group ($\Lambda^0_{t,g}(\omega^S)$).

3. Solve the life-cycle problem assuming that agents expectations over the future are given by this initial guess. The life-cycle problem can be solved recursively starting in the last period.

4. Compute the actual distribution of potential partners in each period. These can be computed using the flows between all states. For example, one can start in period 1, where the distribution of individuals is given by the initial conditions, and consider all states (using a grid for the Pareto weight) and all potential transitions into the next period, using the optimal decision rule from the model, and similarly proceed for future periods. This avoids the need for simulation and the resulting simulation noise, which is helpful to increase convergence speed. The result from this step are the implied distributions of singles at each age ($\Lambda^*_{t,g}(\omega^S)$).

5. Update the guess of the distributions using a weighting factor: $\Lambda^1_{t,g}(\omega^S) = \alpha\Lambda^*_{t,g}(\omega^S) + (1-\alpha)\Lambda^0_{t,g}(\omega^S)$

6. Iterate until convergence.

7. After convergence, simulate from the model in order to be able to flexibly compute various statistics.

The model is implemented in Python using Numba and estimated on a HPC cluster.

Note on interpolating the Nash bargaining solution. Recall that the Nash bargaining solution is given by (omitting some indices and the dependence on state variables):

$$\lambda = \text{argmax}_\lambda \left(V_f^C(\lambda) - V_f^S \right) \left(V_m^C(\lambda) - V_m^S \right)$$

During the solution of the model, the value functions in marriage are known only on a grid for the Pareto weight $\{\lambda_1, ..., \lambda_{n_\lambda}\}$, i.e. only the values $\{V_g^C(\lambda_0), ..., V_g^C(\lambda_{n_\lambda})\}$ are known for each partner g. To allow λ to be a continuous state variable, the solution to the bargaining problem must be interpolated in a way that makes it possible for values off the grid to be chosen. It turns out to be most convenient to work with the first-order conditions from the Nash bargaining solution, which can be interpolated smoothly in λ to find a root:

$$\frac{\partial V_m^C}{\partial \lambda}(\lambda)(V_f^C(\lambda) - V_f^S) + \frac{\partial V_f^C}{\partial \lambda}(\lambda)(V_m^C(\lambda) - V_m^S) = 0$$

This equation uses the derivatives of the utility of couples ($\frac{\partial V_g^C}{\partial \lambda}$). These can be estimated via the following formulas. This approximation is well-known through the shape-preserving spline interpolation in Matlab (*PCHIP*) which also uses it. Given a grid $\{(x_i, y_i)\}$, the slope approximation s_i at grid point i is:

$$\begin{aligned} h_i &= x_i - x_{i-1} \\ d_i &= y_i - y_{i-1} \\ w_{1,i} &= 2h_{i+1} + h_i \\ w_{2,i} &= h_{i+1} + 2h_i \\ s_i &= \frac{w_{1,i} + w_{2,i+1}}{\frac{w_{1,i}}{d_i} + \frac{w_{2,i}}{d_{i+1}}} \end{aligned}$$

The slopes at the endpoints are obtained through linear extrapolation:

$$\begin{aligned} s_0 &= s_1 + \frac{s_1 - s_2}{x_1 - x_2}(x_0 - x_1) \\ s_n &= s_{n-1} + \frac{s_{n-1} - s_{n-2}}{x_{n-1} - x_{n-2}}(x_n - x_{n-1}) \end{aligned}$$

Appendix C: Further Data Analysis

In this appendix, I look at the role of observable characteristics in explaining the variability of choices among singles. Table 19 reports the mean absolute deviation (MAD) of the time use choices with different sets of controls. In these cases, the MAD is computed based on the residuals from a regression of the time use choice on a series of dummy variables. For example, column (2) includes dummies for the number of kids. Column (3) further includes age bins (with 5 years for each bin) and column (4) includes education dummies (there are three categories in the data: below secondary/seconday/above secondary). To make the interpretation easier, rows 4-6 also show the change in the MAD relative to the case in no controls in percentage terms. Controlling for the presence of children changes the MAD of work and leisure relatively little, but reduces the variability of home hours by 18%. Further including age and education among the control variables somewhat reduces variability, although by not much: the variability of work hours is reduced by at most 12%, the MAD of leisure by 2$ and the MAD of home hours by 0.25%. Overall, therefore, the table suggests that the observed variation in time use choices cannot be easily be explained by observable variables, emphasising the need for unobserved heterogeneity.

TABLE 19: Mean absolute deviation (MAD) with different sets of controls

	No controls	+ Number of Kids	+ Age Bins	+ Education
Work	11.0	10.79	10.13	9.69
Leisure	8.0	8.05	7.75	7.84
Home Hours	8.5	7.0	6.66	6.4
Work (Change in %)	-	-0.02	-0.08	-0.12
Leisure (Ch. in %)	-	0.01	-0.03	-0.02
Home Hrs (Ch. in %)	-	-0.18	-0.22	-0.25

Appendix D: Data requirements for backing out singles' preferences directly

In this appendix, I discuss the mapping between allocations and singles' preferences in this model and under which circumstances it would be possible to back out preferences directly using analytic formulas derived from the first-order conditions. For example, consider a simplified model in which singles consume two goods (private and public) and in which time use is exogenous. If one were to assume Cobb-Douglas preferences over these two goods, it would be possible to estimate singles' preferences directly *on the individual level* by using the budget shares of private and public goods as an approximation for the Cobb-Douglas coefficients. This raises the question whether it is possible to do something similar in the context of my model, for example by converting the fraction of total time spent on work, leisure and chores into preference coefficients. Such a calculation would be a simple and transparent way of interpreting the data in terms of preference heterogeneity.

Recall that time use consists of work hours (h), leisure (l) and home hours (d) and consumption consists of private (c) and public (C) consumption, and that the preference specification is given by:

$$u(c,l,D) = \alpha_c \frac{c^{1-\gamma}}{1-\gamma} + \alpha_l \frac{l^{1-\gamma}}{1-\gamma} + \alpha_D \frac{D^{1-\gamma}}{1-\gamma}$$

The problem of a single agent is:

$$\begin{aligned} \max u(c,l,D) \text{ s.t. } c + C &= wh \\ h + l + d &= 1 \\ D &= F(C,d) \end{aligned}$$

The Lagrangian is:

$$L = u(c,l,D) - \lambda(c + C - wh) - \mu(h + d + l - 1)$$

The first-order conditions are:

$$\begin{aligned} \frac{\partial L}{\partial c} &= \frac{\partial u}{\partial c} - \lambda = 0 \iff \alpha_c c^{-\gamma} = \lambda \\ \frac{\partial L}{\partial l} &= \frac{\partial u}{\partial l} - \mu = 0 \iff \alpha_l l^{-\gamma} = \lambda \\ \frac{\partial L}{\partial h} &= w\lambda - \mu = 0 \iff w\lambda = \mu \\ \frac{\partial L}{\partial d} &= \frac{\partial u}{\partial D}\frac{\partial D}{\partial d} - \mu = 0 \iff \alpha_D D^{-\gamma}\frac{\partial D}{\partial d} = \mu \\ \frac{\partial L}{\partial C} &= \frac{\partial u}{\partial D}\frac{\partial D}{\partial C} - \lambda = 0 \iff \alpha_D D^{-\gamma}\frac{\partial D}{\partial C} = \lambda \end{aligned}$$

The question is whether the preference coefficients $(\alpha^c, \alpha_l, \alpha^D)$ can be expressed as a function of variables which are in principle observable. Note that in the UKTUS data, only time use choices (l, d, h) are observed. In addition, one could also use the Family Resources Survey for the UK for information on consumption, separating aggregate private and public consumption like Lise and Seitz (2011). A limitation of the UK data is that only (1) time use and consumption data are not available in the same data set and (2) only aggregate private consumption $(c_f + c_m)$ can be computed.

As shown by the first-order conditions, directly computing preference coefficients from the data would require more information than available in the UK. With more data, the preference coefficients could at least partly be backed out. For example, in a dataset where consumption and leisure are observed, the ratio between α_c and α_l could directly be computed from the data given this preference specification by combining the first three FOCs:

$$\frac{\alpha_c}{\alpha_l} = \left(\frac{l}{c}\right)^{-\gamma} \frac{1}{w}$$

Backing out α_D is more difficult as it requires computing D, which depends on the parameters of the home production function. Even conditional on the estimated home production parameters, it would require simultaneously observing time use (l, d) and public consumption (C) on the individual level. In the UK context, this means that directly backing out preferences from the data is not possible. Instead, the optimisation problem must be solved numerically. However, it would be interesting to explore such a direct mapping between allocations and preferences in a richer dataset.[35]

[35]For example, the Dutch LISS panel or the Japanese data used in Lise and Yamada (2019) allow to simultaneously observe consumption, time use and wages.

Appendix E: Additional Tables and Figures

TABLE 20: Sensitivity and specificity for utility-poverty

True Welfare Measure Chosen Welfare Measure	Sensitivity MMWI	Specificity MMWI
(1) Utility	0.76	0.94
(2) Marginal utility	0.77	0.94

Notes: This table shows sensitivity and specificity for identifying welfare-poor individuals for utility-based welfare measures. In the first row, the individuals with the lowest utility are classified as poor. In the second row, the individuals with the highest marginal utility (with respect to a small increase in non-labour income of the household) are chosen.

TABLE 21: Rank correlations between welfare measures

	Household income	Individual consumption	Empirical sharing rule
Rank correlation (MMWI)	0.72	0.75	0.87
Rank correlation (Rente Criterion)	0.84	0.93	0.73
Rank deviation (MMWI)	16.70	15.60	10.00
Rank deviation (Rente Criterion)	12.80	7.50	16.30

Notes: This table show the rank correlations (using Spearman's rho) between different welfare measures and the Rente Criterion and MMWI (rows 1 and 2). In addition, rows 3 and 4 show the mean absolute deviation in ranks. This can be interpreted as the mean error one would make when for instance predicting the welfare rank of an individual based on household income. The rank deviation reflects the underlying scaling of ranks between 0 (lowest) and 100 (highest).

TABLE 22: Rank correlations between partners

	Rank correlation	Mean absolute rank deviation
Household income	1.00	0.00
Individual consumption	0.72	16.35
Empirical sharing rule	0.84	11.93
MMWI	0.82	12.86
Rente Criterion	0.56	21.45

Notes: This table shows the rank correlations (using Spearman's rho) between partners for each of the welfare measures. In addition, the table also reports the mean absolute deviation of the ranks of each partner (also see notes of Figure 21).

FIGURE 10: Distribution of Pareto weights when there is no intra-household inequality

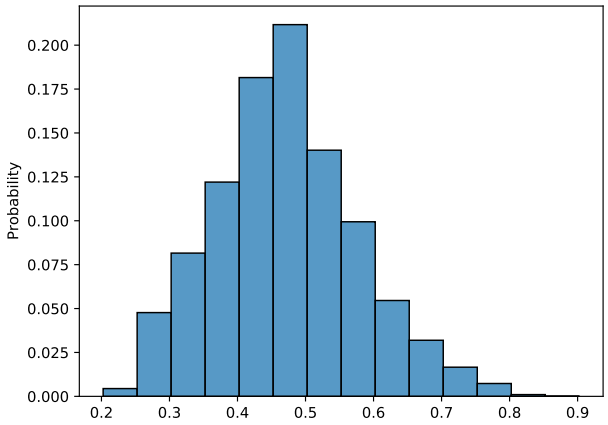

Notes: This figure shows the distribution of hypothetical Pareto weights when intra-household inequality is eliminated (i.e. Pareto weights are adjusted so that the MMWI of each person is equal).

FIGURE 11: Impact of Direct and Reoptimisation Effect along the Income Distribution

(a) Direct Effect

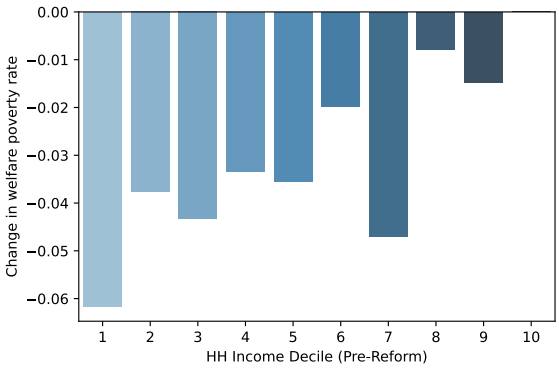

(b) Reoptimisation Effect

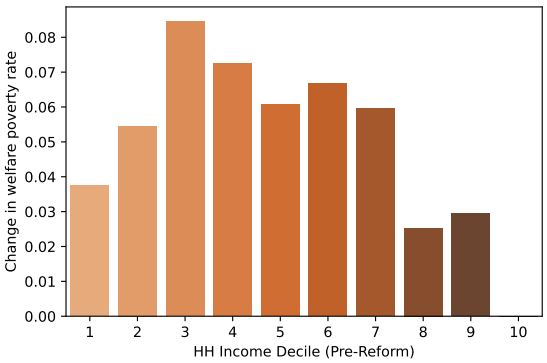

Notes: This figure shows the impact of the direct and the reopmisation effect on welfare poverty rates according to the MMWI (see main text for definition of the effects). On the x-axis, there is the pre-reform household income decile of the individual. Note that the income decile of an individual can change due to the reform.