



Centre for
Economic
Performance

Discussion Paper

ISSN 2042-2695

No. 1954
October 2023

**Individual
welfare
analysis:
A tale of
consumption,
time use and
preference
heterogeneity**

Tim Obermeier



THE LONDON SCHOOL
OF ECONOMICS AND
POLITICAL SCIENCE ■



Economic
and Social
Research Council

Abstract

How accurately does household income reflect the well-being of the individuals living within the household? Looking at household income does not take unequal consumption sharing within families, the value of time use (leisure and housework) and preference heterogeneity into account. I build a model of family decision-making and the marriage market which jointly captures these aspects and estimate the model based on British time use data. I use the estimated model to study poverty and inequality based on the individual-level Money-Metric Welfare Index (MMWI). The main result is that only 59% of individuals who are poor in terms of the MMWI ('welfare-poor') are also income-poor, suggesting that the conventional focus on income misses a substantial fraction of the welfare-poor. I find that accounting for unobserved preference heterogeneity is an important factor in assessing individual welfare. From an aggregate perspective, inequality within families accounts for 18% of overall welfare inequality, and heterogeneity in economies of scale across households account for 23% of welfare inequality. 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 Codes: E21; I32; D13; D63

This paper was produced as part of the Centre's Growth Programme. The Centre for Economic Performance is financed by the Economic and Social Research Council.

I'm grateful to Richard Blundell, Christian Bredemeier, Pierre-André Chiappori, Sena Coskun, Monica Costa Dias, Bram De Rock, Falko Jüßen, Jan Gravert, Paula Gobbi, Anne Hannusch, Karen Kopecky, Valérie Lechene, Krishna Pendakur, Franck Portier, Michèle Tertilt and John Van Reenen, as well as participants at the Workshop in Gender and Family Economics (Cergy-Pontoise), the 2022 annual conference of the EEA (Milan), the 2022 annual conference of the VfS (Basel) as well as seminar participants at IFS, UCL, the University of Wuppertal and various other universities for many very helpful comments. The paper was previously circulated as Individual Welfare Analysis: What's the Role of Intra-Family Preference Heterogeneity?. Support by the state of Baden-Württemberg through bwHPC is gratefully acknowledged.

Tim Obermeier, University of Leicester, IFS and Centre for Economic Performance at London School of Economics.

Published by

Centre for Economic Performance

London School of Economic and Political Science

Houghton Street

London WC2A 2AE

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means without the prior permission in writing of the publisher nor be issued to the public or circulated in any form other than that in which it is published.

Requests for permission to reproduce any article or part of the Working Paper should be sent to the editor at the above address.

© T. Obermeier, submitted 2023.

1 Introduction

The discourse surrounding inequality often focuses on household income and a vast literature has studied household income inequality and its implications (see e.g. [Atkinson and Bourguignon \(2014\)](#)). Let's consider a simple example involving a childless couple consisting of partner A (Alex) and partner B (Bertie), who have a total monthly household income of £2100. The conventional approach to measuring inequality is to adjust household income for economies of scale (the fact that some goods, like housing, are public goods within the household) and, using the OECD equivalence scale, assign a value of £1400 ($=£2100/1.4$) to each person. How accurately does this calculation reflect the economic well-being of Alex and Bertie? The answer to this question is important for the analysis of inequality and poverty in labour and macroeconomics.

A substantial body of work in the literature on collective household models has pointed out that there are several reasons why household income is not a perfect measure of individual well-being (see, for instance, [Chiappori and Meghir \(2015\)](#), [Chiappori and Mazzocco \(2017\)](#), [Lise and Seitz \(2011\)](#), [Cherchye et al. \(2015\)](#), [Cherchye et al. \(2018\)](#), [Cherchye et al. \(2018\)](#), [Cherchye et al. \(2020\)](#)). Consumption can be shared unequally within families, so that one partner can have access to a larger fraction of household resources. In addition, individuals also value leisure and domestic work in addition to consumption. Finally, the mapping between income and well-being also depends on individual preferences and how each household member values consumption, leisure and domestic activities. For example, [Lise and Seitz \(2011\)](#) estimate that 25% of consumption inequality among childless British households is due to inequality *within* households. This suggests that the simple calculation above, which assigns £1400 to both partners, does not give a full picture of inequality as it assumes an equal division within the household.

Among these aspects, the role of preference heterogeneity has attracted least attention. To see why preference heterogeneity is important, suppose that Alex and Bertie consume two goods, a private and a public good, and spend their time on market work, leisure and housework. Alex spends £400 on private consumption and Bertie spends £700. In addition, Alex's contributes more in terms of housework and has only 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 key point is that welfare conclusions about who is better off depend strongly on the preferences of Alex and Bertie. With equal preference, Bertie is better off as they consume more and has more leisure, and both partners have the same level of the public good. With preference heterogeneity, the conclusion could be the opposite: Alex' lower level of consumption does not necessarily reflect that they are in a disadvantaged position within the household. Instead, an equally plausible scenario is that their priorities are different. Alex could care less about private consumption, such as clothing, and derive more utility from the household-level public good. As a result, Alex could be better off than Bertie or there could be no inequality between the two at all.

In this paper, I study the relationship between household-level income and individual-level well-being in a model that *jointly* accounts for unequal consumption, time allocations *and* preference heterogeneity. The model is based on [Lise and Yamada \(2019\)](#) and considers a setting with three goods - private consumption, leisure and a public good -, but introduces preference heterogeneity and embeds the household decision problem in a more structural framework with a marriage market, following the literature on structural models of family behaviour (e.g. [Shephard \(2019\)](#), [Low et al. \(2021\)](#), [Bronson and Mazzocco \(2019\)](#)).¹ 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. In addition, couple formation and dissolution are endogenous through a search-based marriage market. The key novelty in this paper is that I introduce population-wide preference heterogeneity, which is *unobserved* to the econometrician. This is an important distinction to previous studies, which only allow for limited preference heterogeneity by linking preferences to observable characteristics, and leads to a more realistic extent of heterogeneity.²

The model can be estimated based on time use datasets, complemented with information on wages, both of which are commonly available. 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), while USOC contains information on wages. The key empirical challenge is that individual preferences are unobserved in the data, as there are no direct measures of the preferences needed in collective household models. Unobserved preference heterogeneity is difficult to incorporate into microeconomic estimation strategies, such as the GMM strategy in [Lise and Yamada \(2019\)](#), since preference heterogeneity needs to be separated from unobserved differences in the extent of sharing within the household.³ Because of this issue, the more structural approach with the marriage market equilibrium is needed. The 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 *joint* distribution of partners is determined endogenously along with the distribution of singles. This follows the approach of ‘identification from market equilibrium’ ([Chiappori and Meghir \(2015\)](#)), which has

¹The approach of [Lise and Yamada \(2019\)](#) is model the household decision problem, but not the marriage market, and estimate the model by using GMM based on the first-order conditions of the decision problem.

²For example, a model with observables can allow for preference heterogeneity between men and women, but it is a strong assumption to say that *all* women care more (or less) about a good than *all* men; instead, there is likely to be a lot of variation within the group of men and women (which is unrelated to other observables, such as education). Unobserved preference heterogeneity is more realistic because it better captures the within-cell variation in preferences.

³As an example, suppose we observe that a woman consumes more than her partner. This could either be 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 is difficult in the GMM strategy of [Lise and Yamada \(2019\)](#), as it requires disentangling the two, and requires modelling the marriage market.

also been applied in related contexts.⁴

The analysis leads to four main messages. First, I use the estimated model for an individual-level poverty analysis. Studying poverty has been a common topic in the collective model literature (e.g. Cherchye et al. (2015), Cherchye et al. (2018)). As a welfare measure in collective household models, Chiappori and Meghir (2015) have proposed the *Money-Metric Welfare Index* (MMWI). This converts an individual's utility from consumption, leisure and the home good into a monetary index that can be compared across individuals with heterogeneous preferences. I classify individuals as income-poor if they are in the bottom quintile of the (household-) income distribution, and as MMWI-poor (or 'welfare-poor') if they are in bottom quintile of the distribution of the MMWI. I find that only 59% of welfare-poor individuals are also income-poor, meaning that a substantial fraction of those who are least well-off in terms of welfare live *outside* of income-poor households.

Second, and relatedly, the analysis shows the relative importance of the three factors that drive a wedge between income and the MMWI - (1) consumption sharing within the family, (2) time use and (3) preference heterogeneity. For example, if the differences between household income and the MMWI are mainly driven by consumption inequality within families, consumption-poverty should capture much better who is welfare-poor than income-poverty. However, looking at consumption only identifies 56% of welfare-poor individuals, showing that consumption sharing is not the main factor. An index that combines individual consumption and the market value of leisure and home production does slightly better than income and identifies 64% of welfare-poor individuals, but still leaves 36% of individuals unidentified. This means that it is difficult to assess who is welfare-poor without explicitly taking information about preferences into account. In other words, the *interaction* between consumption sharing, time use and preferences is key for individual welfare and information on all of these three aspects is needed to assess an individual's position in the MMWI distribution.

Third, I assess the role of intra-household inequality and economies of scale for welfare inequality more broadly. While the poverty analysis is closely related to the collective model literature, from a macroeconomic perspective it is also interesting to look at overall inequality. In a standard variance decomposition, intra-household inequality accounts for 18.1% of overall welfare inequality (that is, inequality in terms of the MMWI). This suggests that inequality *within* households, which is often abstracted from, is relevant for overall inequality. Inequality in economies of scale account for 23.1% of welfare inequality. Economies of scale are heterogeneous across households in the model because of preference heterogeneity: couples where both partners have a high preference for the public good can share a lot of resources and enjoy high economies of scale. Taken

⁴For 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. (2021)). 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.

together, these results suggest that intra-household inequality and economies of scale together make a significant contribution to aggregate welfare inequality.

Finally, I illustrate the relevance of using individual-level welfare measures for policy analysis, in the context of minimum wage policy. A well-known fact is that many minimum wage earners do not live in income-poor households but are often located in the middle of the income distribution, suggesting that minimum wages are not ideally targeted for reducing poverty. I use the estimated model to study how a hypothetical minimum wage increase of 20% affects welfare-poverty from an individual perspective. The main finding is that there is a substantial reduction in welfare-poverty particularly among minimum wage earners whose partner has an hourly wage between £20-25, where the household is less likely to be considered as income-poor in the traditional sense. Studying the role of the mechanisms behind this reduction in welfare-poverty, I find that the bargaining effect, which is only present in collective household models, plays a key role for this result: increasing the minimum wage increases the intra-household decision power of minimum wage earners and thereby their welfare. This finding suggests that it is important to study minimum wage policy from the perspective of a collective rather than a unitary household model.

Related literature and contribution. First, the main contribution of the paper relates to the growing literature on measuring consumption and welfare on the level of individuals.⁵ Chiappori and Meghir (2015) and Chiappori and Meghir (2014) theoretically discuss how to measure individual welfare in collective household models and propose the *Money-Metric Welfare Index* (MMWI). However, there has been very little empirical or quantitative research of poverty and inequality in terms of the MMWI. The large literature on individual-level welfare in collective models typically focuses on consumption, both in developed (e.g. Cherchye et al. (2020), Cherchye et al. (2015)) and developing countries (e.g. Bargain, Donni, and Kwenda (2014), Calvi (2020), Penglase (2021), Brown, Calvi, and Penglase (2021), Tommasi (2019)). So far, only Cherchye et al. (2018) study poverty in terms of the MMWI. They use US data from the PSID and build a model of consumption and leisure with limited preference heterogeneity, where preferences are allowed to differ only between men and women but not within gender. The main contributions of my paper are to study individual poverty in terms of the MMWI while simultaneously accounting for consumption sharing, time use and rich unobserved preference heterogeneity, and to quantify the role of intra-household inequality and economies of scale for the wider distribution (overall welfare inequality in terms of the MMWI). The analysis of aggregate inequality also relates to the macroeconomic literature, which has mostly focused on income, wage or consumption inequality, and illustrates the role of the MMWI for the analysis of overall welfare inequality.⁶

⁵Various 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)).

⁶Attanasio and Pistaferri (2016) point out that a welfare analysis should ultimately take the value of different consumption goods and leisure time into account. Boerma and Karabarbounis (2020) study welfare inequality while taking the value of non-market work into account. Fang, Hannusch, and Silos (2021) develop a model with different

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. \(2021\)](#), [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 the MMWI and the endogenous allocation of consumption and time.

Finally, in terms of modelling, the paper contributes 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 \(2012\)](#)). This can be restrictive in some 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). However, *unobserved* (individual-level) preference heterogeneity is challenging to incorporate because of the identification issues discussed earlier and has mostly been abstracted from. A notable exception is [Cherchye et al. \(2020\)](#), who also use marriage market restrictions for identification, but in a very different framework as they follow a revealed preferences approach with a static frictionless marriage market. In addition, [Balleer, Merz, and Papp \(2021\)](#) follow a different approach by estimating preference heterogeneity based on a non-cooperative model of household behaviour. My paper complements previous work by adding unobserved preference heterogeneity to a fully parametric structural collective model following the search literature (e.g. [Shimer and Smith \(2000\)](#), [Goussé, Jacquemet, and Robin \(2017\)](#), [Mazzocco, Ruiz, and Yamaguchi \(2013\)](#), [Shephard \(2019\)](#), [Ciscato \(2019\)](#), [Obermeier \(2019\)](#)). Compared to revealed preference models, which have the advantage of being nonparametrically identified, this approach comes with a different set of strengths. [Cherchye et al. \(2020\)](#) analyse well-being in terms of expenditure, using a concept called *relative individual cost of equivalent bundle* (RICEB), whereas the model from this paper allows to compute welfare/the MMWI. In addition, the parametric search approach allows to conduct counterfactual experiments and is therefore a useful complement to revealed preference models, for example to study the impact of policies on individual welfare like in the minimum wage analysis presented here.⁷

The rest of this paper is organised as follows. Section 2 describes the model, which is estimated in section 3. Section 4 contains the welfare and policy analysis. Section 5 concludes.

kinds of consumption and leisure goods (luxuries and necessities) and analyse welfare inequality.

⁷In this paper, the marriage market outcomes and the sharing rule (Pareto weights) are endogenous, which allows to study how they are affected by policy changes (or other counterfactual experiments). While [Cherchye et al. \(2020\)](#) identify the sharing rule from the data, they do not model how it is determined, which is important for counterfactual experiments.

2 Model

2.1 Motivation and Overview

The purpose of the model is to study how household income is related to the welfare of the individuals living in the household. To address this question, having a structural model is essential for multiple reasons. First, studying welfare requires an explicit model of utility. Second, as is common in the collective model literature, not everything is directly observed in the data (see e.g. [Chiappori and Meghir \(2015\)](#)). Both the consumption allocation within families as well as individual preferences are backed out indirectly through the combination of the available data and the structure of the model.

The model is closely related to the setting from [Lise and Yamada \(2019\)](#) and embeds a very similar household decision problem into a more structural framework following the literature on structural models of household decision-making and the marriage market (also e.g. [Goussé, Jacquemet, and Robin \(2017\)](#), [Mazzocco, Ruiz, and Yamaguchi \(2013\)](#), [Low et al. \(2021\)](#), [Bronson and Mazzocco \(2019\)](#), [Ciscato \(2019\)](#), [Shephard \(2019\)](#)). The household decision problem is a setting with three goods: personal consumption (c_i), leisure (l_i) and a public 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.⁸ Linking the household decision problem to a simple dynamic marriage market endogenises the allocation of decision power within the couple and determines the matching of the unobserved preference types.

2.2 Demographics

The model contains a two-sided matching market with two populations, which will correspond to men and women in the empirical analysis.⁹ 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 .¹⁰ The 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. Having children is a binary state ($b \in \{0, 1\}$), which can be thought of as each couple having two children.

⁸Note 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.

⁹Note 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.

¹⁰As 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.

The arrival of children is determined by an exogenous probability p_b and children grow up and leave the household with probability p_g . 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. $\alpha_i^c = 1 - \alpha_i^D - \alpha_i^l$).¹¹

The preference coefficients determine the weight of each good in the utility function, where γ is a common curvature parameter and η governs the substitutability between the goods:

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

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 (α_i^D). 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.

In terms of domestic production, each partner in a couple contributes a time input (d_f and d_m for women and men). These time inputs are aggregated into a total domestic time input for the couple:

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

$\theta_g(b)$ is the relative 'productivity' of time inputs, which could also reflect for instance social norms relating to the division of labour within couples. This parameter is allowed to depend on

¹¹This can be thought of as a *no-utility-monster restriction*, as it rules out that individuals obtain more utility from each good.

the presence of children. z determines how easily male and female home hours can be substituted. ζ_f and ζ_m are additional productivity scales which will govern intra-household division of market work. For singles, the total domestic time input is given by:

$$\bar{d} = \theta(b)d$$

The public good is produced using both the aggregate time input (\bar{d}) and public expenditure of the household (q):

$$D(\bar{d}, q) = A(b)(\omega_H \bar{d}^{z_2} + (1 - \omega_H)q^{z_2})^{\frac{1}{z_2}}$$

ω_H is the weight on time inputs relative to money inputs and z_2 determines the substitutability between time and money. $A(b)$ is a TFP parameter which allows home time to be more productive when children are in the household, to capture the impact of children on family time allocations.

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

$$\begin{aligned} \max_{c, q, l, d, h} \quad & u_i(c, l, D) \\ & c + q = \bar{w}_i h \\ & l + d + h = 1 \\ & D = D(q, d) \end{aligned}$$

The time budget of the individual is normalised to 1. The wage \bar{w}_i can therefore also be interpreted as the full income of the individual. 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):

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

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} :

$$\begin{aligned} \max_{c_i, c_j, q, l_i, l_j, d_i, d_j, h_i, h_j} \quad & \lambda_{ht} u_i(c_i, l_i, D) + (1 - \lambda_{ht}) u_j(c_j, l_j, D) \\ c_i + c_j + q \quad &= \bar{w}_i h_i + \bar{w}_j h_j \\ l_i + h_i + d_i \quad &= 1 \\ l_j + h_j + d_j \quad &= 1 \\ D \quad &= D(q, d_i, d_j) \end{aligned}$$

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.¹² 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.¹³

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:

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

¹²When partners have equal preferences, $\lambda_{ht} = 0.5$ corresponds to equality within the household. With unequal preferences, the Pareto weight that induces equality typically deviates from 0.5.

¹³See Blundell, Chiappori, and Meghir (2005) for a theoretical discussion of the conditions required for this effect.

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 market and lead to sorting.¹⁴

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 ($\theta \sim N(0, \sigma_\theta)$). 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_{pt}^S 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_{ht}^C (which includes the Pareto weight) and each partner p obtains utility $V_{pt}^C(\omega_{ht}^C)$, where V_{pt}^C 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 each partner $p \in \{i, j\}$:

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

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_{ht}^C):

$$\tilde{\lambda} = \operatorname{argmax}_\lambda \left(V_{it}^C(\omega_{ht}^C) - V_{it}^S(\omega_{it}^S) \right) \cdot \left(V_{jt}^C(\omega_{ht}^C) - V_{jt}^S(\omega_{jt}^S) \right)$$

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 \(2017\)](#).¹⁵

¹⁴In practice, κ is identified as the residual required to match the extent of wage sorting observed in the data.

¹⁵Assuming 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.

Given these marriage market decisions, the continuation value of singlehood can be expressed more explicitly as the expectation over the marriage market outcomes in the next period:

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

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 $V_{i,t+1}^S(\omega_{it}^S)$ 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).¹⁶ 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.

Models with separation require some assumptions about how children are dealt with in separation. As a simple way of dealing with children, which avoids additional computational complexities, I assume that both men and women remain in the child state ($b = 1$) after separation, which can be thought of as capturing that both partners contribute to the child without explicitly modelling the interaction between them. In addition, I assume that children grow up upon remarriage, capturing that remarriage typically occurs relatively later in life. Note that since the focus of the model is on the cross-section and there are no savings or human capital accumulation decisions, this simplification is unlikely to influence the results from the analysis.¹⁷

¹⁶For 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.

¹⁷More complex ways of modelling custody are possible, but would increase the computational burden in solving the model.

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

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

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.

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

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 proposed the *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 $\bar{u} = u_i(c_i, l_i, D)$ 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 (\bar{w}). In order to compute the MMWI, one first needs to solve for the indirect utility function of the single problem:

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

Note that u_i depends on the preferences of the individual. Then, the MMWI (denoted as M_i) is implicitly defined by:

$$V_i(M_i) = \bar{u} \Leftrightarrow M_i = V_i^{-1}(\bar{u})$$

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. When looking at singles, the MMWI of a single is by definition given by their full income \bar{w}_i (as it is defined as the full income someone would need in singlehood to achieve a given utility level). Since full income can equivalently be expressed as the hourly wage rate, this amounts to measuring inequality in terms of wages rather than income. For coupled individuals, the MMWI further takes economies of scales and intra-household inequality into account, both of which will be discussed in detail in the analysis.

Note that there are multiple ways of comparing welfare in the presence of preference heterogeneity, as has been pointed out by the literature in welfare economics (see e.g. [Fleurbaey \(2006\)](#), [Decoster and Haan \(2015\)](#)). In Appendix A, I discuss in more detail how the MMWI relates to the welfare criteria introduced by [Fleurbaey \(2006\)](#) and how the main results would be affected by using an alternative welfare criterion. From the perspective of the welfare measures defined in [Fleurbaey \(2006\)](#), the MMWI is very similar to the "Wage Rate Criterion", which compares individuals according to their wage. The idea behind this is easiest to illustrate in the context of two singles. Suppose both singles earn the same wage, but one of them works part-time while

the other works full-time due to preference heterogeneity. An income-based welfare measure would say that the individual who works part-time is worse off due to their lower labour earnings. According to the Wage Rate Criterion/MMWI, by contrast, both are equally well off. The Wage Rate Criterion seeks to eliminate the part of inequality which is due to choices (work hours) and focuses on the inequality which is due to exogenous constraints (wages). [Fleurbaey \(2006\)](#) also introduces an additional welfare measure, the "Rente Criterion", which also takes inequality due to choices into account. In the main analysis, I use the MMWI as this follows the previous literature on intra-household inequality ([Cherchye et al. \(2018\)](#), [Chiappori and Meghir \(2015\)](#)).

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 1](#) shows the summary statistics for the sample.

TABLE 1: 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 on separation and marriage rates.

3.2 Model specification

In practice, the time horizon is set to $T = 20$. 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. The preference type of each individual is described by a vector $(\alpha_i^C, \alpha_i^L, \alpha_i^D)$. Therefore, we need a preference distribution $F(\alpha^C, \alpha^L, \alpha^D)$, which describes how the population is divided into different preference types. The main issue that needs to be dealt with is that the distribution has to fulfil the restrictions $\alpha_i^C + \alpha_i^L + \alpha_i^D = 1$ and $0 \leq \alpha_i^K \leq 1$ for $K \in \{C, L, D\}$. This essentially requires specifying a probability distribution on the two-dimensional triangular space shown in Appendix Figure 6, which needs to be discrete for computational reasons.

In practice, I draw auxiliary normal random variables $\chi^K \sim N(\mu_K, \sigma_K)$ ($K \in \{C, L, D\}$) with mean μ_K and variance σ_K . Then the preference coefficients are defined as:

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

As a parsimonious way of discretising this distribution, I discretise at the level of α^K rather than χ^K .¹⁸ Given the parameters of the distribution, I simulate a sample of (α_i^C, α_i^L) . Then, I group the simulated sample into n^D equally sized bins according to the distribution of α_i^C . Within each bin, I further create n^D additional equally sized bins according to α_i^L . This, overall, results in $n^P = n^D \cdot n^D$ bins and for each of these I compute the mean values of α_i^C and α_i^L as grid points. Finally, the each point from the simulated sample is assigned to the closest grid point to compute the probabilities of each point. Taken together, the preference distribution is parametrised by five parameters $(\mu^C, \mu^L, \sigma^C, \sigma^L, \sigma^D)$, which through the steps described here lead to a discrete distribution $\{\alpha_k^C, \alpha_k^L\}_{k=1, \dots, n^P}$, which is an input into the model. In practice, I set $n^D = 5$.

The wage distribution is estimated externally based on the USOC data. In order to generate a wage distribution which is suitable for the minimum wage analysis, I use a particular procedure rather than simply fitting a discretised log normal distribution. I choose 10 discrete points between the minimum wage and an upper bound. The lowest point corresponds to the minimum wage and the second grid point is equal to 1.3 times the minimum wage (this easily allows for a policy experiment of raising the minimum wage while only affecting the lowest grid point). The other points are chosen based on the quantiles of the remaining distribution. Then, the probability mass of each point is computed by assigning each wage from the sample to the closest grid point and computing the fractions.

The final part of the model specification is that I parametrise the additional home productivity

¹⁸Discretisation on the level of χ^K is more strongly affected by the curse of dimensionality, since for example 3 grid points for each random variable would already result in 27 grid points for the preference distribution, whereas discretising α_K requires only discretising in two dimensions and computing the last one via $\alpha_i^C + \alpha_i^L + \alpha_i^D = 1$.

parameters as follows:

$$\zeta_i = \left(\frac{w_i}{\bar{w}_{g(i)}} \right)^\delta$$

w_i is the wage rate of the individual and $w_{g(i)}$ is the average wage of the gender of the individual. This assumption makes high-wage individuals more productive in home production while in couples (which could for example capture social norms relating to the division of labour within households), which increases their contribution to domestic work relative to the case without this assumption. This follows [Alon, Doepke, and Coskun \(2018\)](#) and is required because otherwise it would be more difficult for the model to generate a realistic extent of intra-household specialisation.¹⁹

3.3 Estimation strategy and identification

Some model parameters are set externally. The yearly discount rate β is set 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 remaining parameters are estimated to match a set of data targets, by minimising the distance between model and data moments. Overall, 16 parameters are estimated (see Table 2): those of the preference distribution ($\mu^C, \mu^L, \sigma^C, \sigma^L, \sigma^D$), the substitutability parameter of the utility function (η), the home production TFP parameter for children ($A(b = 1)$), the home productivity parameters ($\theta_{f,b=0}, \theta_{f,b=1}, \omega_H, z, z_2, \delta$) and the parameters for the love process and homophily (ρ_l, σ_l, κ). Table 3 shows the target moments which are used to pin down the parameters. As is usually the case, all parameters jointly determine all moments. However, for different parameters, there are some moments that are particularly informative about them, so that it is possible to provide some intuition for identification by discussing these links.

The mean parameters of the preference distribution, along with the home production parameters, are identified by targeting the mean time use choices. These moments include average work hours and home production hours of men and women (leisure is omitted since it is the residual). In addition, I also include women's average work hours for different bins of the income of their partner (for example, the moment "men's income 3" refers to men earning between £30,000-39,999 and the other bins are similarly defined). These moments make sure that the model generates a realistic amount of intra-household specialisation in terms of how strongly women reduce their work hours when their partner earns more. In the data, women who are in a relationship with high income men work modestly less than those who are not.²⁰ The parameters which have a strong influence on the extent of specialisation are the substitutability between partners' time

¹⁹Without this assumption, a collective model with home production would predict that women's labour supply rapidly goes to zero when the wage rate of their partner rises unless z takes unrealistically low values, making the time inputs into home production very strong complements.

²⁰These moments are computed based on the British Household Panel Survey (Understanding Society/USOC).

TABLE 2: Estimated Parameters

Parameter	Description	Value
μ^L	Preferences, mean parameter 1	1.13
μ^C	Preferences, mean parameter 2	0.39
σ^C	Preferences, variance parameter 1	0.61
σ^L	Preferences, variance parameter 2	0.39
σ^D	Preferences, variance parameter 3	0.85
z	Substitutability, time	-0.49
z_2	Substitutability, money/time	0.50
$\theta_{f,b=0}$	Home productivity, women, no kids	0.73
$\theta_{f,b=1}$	Home productivity, women, with kids	0.77
$A(b = 1)$	Home production TFP (children)	3.79
ω_H	Weight on time, home production	0.74
δ	Home productivity scale	1.04
η	Substitutability, preferences	0.14
σ_l	Variance, love shock	1.34
ρ_l	Autocorrelation, love shock	0.93
κ	Homophily parameter	1.41

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

inputs in the home production function (z) and the home production scale (δ). Since some of the parameters depend on the presence of children (the TFP parameter $A(b)$ and relative productivity $\theta_{f,b}$), I match time use choices separately by whether the couple has children.

The variance parameters of the preference distribution ($\sigma^C, \sigma^L, \sigma^D$) target the variability of time use choices of singles. For example, the more dispersed the distribution of work hours among singles, the higher the variation in consumption preferences (similarly for hours spent on leisure and home production and the corresponding preferences). Here, the marriage market is crucial for identification as the marriage market equilibrium determines who stays single and who marries whom. Conditional on the parameters, the model endogenously generates a distribution of singles, as well as the joint distribution of partners' preferences, while taking selection into singlehood/relationships based on preferences and wages into account. The larger the variance parameters, the larger the variability of time use choices among singles. As a measure of variability, I choose the interquartile range (the difference between the 75th and 25th percentile), which is more robust to outliers than for example the variance, and divide it by the median as a standardisation. The larger the variance parameters, the higher the model-implied interquartile range and so the variance parameters can be chosen to match the interquartile range from the data. Note that an important assumption is that marriage per se does not change preferences, as otherwise the link between the preference distributions of couples and singles would be more complicated. The assumption of preference stability in marriage has also previously been used in the previous literature on collective models (based on [Browning, Chiappori, and Lewbel \(2013\)](#)). Also 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 1](#).

Finally, the last set of parameters and moments relates to the marriage market. The parameters are the autocorrelation of the love shock (ρ_l), the variance of the love shock (σ_l) and the homophily parameter (κ). The autocorrelation and the variance of the love process determine the incentives for searching for partners on the marriage market and to separate if love deteriorates. For example, if the autocorrelation and the variance is high, individuals wait longer to find an acceptable match. The autocorrelation also has a strong influence on the separation rate: if the love shock is very persistent there will be fewer separations over time. These two parameters are estimated to match the fraction of individuals living in couples and the separation rate. The homophily parameter determines how assortative matching is on wages. The corresponding data target is the wage correlation between partners.

In practical terms, the estimation of the model is computationally fairly expensive since there are 16 parameters to be optimised over and solving the model requires finding a fixed point in each iteration of the optimisation algorithm. I estimate the model on a HPC cluster with 40 cores using Powell's optimisation algorithm which can deal with complex and potentially non-differentiable objective functions. To reduce the estimation time, I estimate the model in a somewhat reduced

specification with fewer wage and preference types.²¹

The fit of the model is good, as shown by Table 3: the model closely matches the time use choices and the key patterns in the data, such as women providing more home production hours than men or children increasing the amount of home hours. The model also matches the variability of singles' time use choices well, as well as the marriage and separation rate, and the relationship between female work hours and men's income. The model underpredicts the wage correlation between partners.

3.4 Model implications: Consumption and marriage market outcomes

Before turning to the welfare analysis, it is useful to first discuss some implications of the estimated model in terms of consumption, assortative matching and selection into marriage. Regarding consumption, note that an advantage of the structural model is that it can be estimated based on time use data, which is useful since detailed data on consumption is often not available. This is similar in closely related models and results from the fact that the bargaining process is modelled via Nash bargaining.²² Therefore, there are no separate parameters for consumption that need to be estimated, but consumption allocations are identified indirectly by estimating the preference and home production parameters together with the assumption on bargaining. However, looking at consumption data is a useful plausibility check to support the modelling assumptions.

Lise and Seitz (2011) use the Family Expenditure Survey (FES) for the UK and categorise expenditure into either private or public. Appendix Figure 7 shows the distribution of the fraction of public consumption relative to total consumption, both for singles and couples, using the data from Lise and Seitz (2011) and their definition of public consumption (housing, heating and household durables). First, note that Lise and Seitz (2011) estimate that childless couples spend around 30% of their expenditure on public goods. The model-implied fraction, while not directly targeted, is very similar (30.6%), so that the public good share in the model is in line with previous literature. This is an important statistic for the extent of intra-household consumption inequality, since a higher public good share reduces the scope for consumption inequality between partners. Second, a key prediction of the model is that preference heterogeneity generates dispersion in the fraction of public expenditure among singles. The data shows that there is indeed a substantial amount of dispersion, which is consistent with preference heterogeneity over private and public

²¹The reduced specification sets the number of wage types to 4 rather than 10 and the number of preference types to 16 instead of 25. The implied moments are mostly very similar when increasing the number of types. The wage correlation between partners and the relationship between female hours and male income are most sensitive to the number of types, which makes it less easy to precisely fit these moments. The moments reported in the fit table correspond to the specification with the full number of types.

²²This differs from papers such as Cherchye et al. (2015) or Cherchye et al. (2018), where consumption data is a key input for the estimation.

TABLE 3: Model Fit

	Data	Model
Work hours, women, mean, couples, no kids	32.08	35.36
Home hours, women, mean, couples, no kids	21.83	21.06
Work hours, women, mean, couples, with kids	22.86	25.18
Home hours, women, mean, couples, with kids	37.35	33.94
Work hours, men, mean, couples, no kids	38.31	44.40
Home hours, men, mean, couples, no kids	10.49	10.79
Work hours, men, mean, couples, with kids	40.33	40.46
Home hours, men, mean, couples, with kids	14.32	16.03
IQR (work)	0.50	0.60
IQR (leisure)	0.64	0.47
IQR (HP)	1.14	1.12
Marriage rate	0.78	0.80
Separation rate	0.04	0.05
Wage correlation	0.27	0.17
Female work hours, mens income 3	26.43	34.45
Female work hours, mens income 4	27.19	30.72
Female work hours, mens income 5	25.90	27.70
Female work hours, mens income 6	24.94	25.91
Female work hours, mens income 7	23.68	23.21
Female work hours, mens income 8	24.36	19.57
Female work hours, mens income 9	21.74	18.91
Female work hours, mens income 10	21.50	20.39

Notes: The table shows the fit of the model by comparing the targeted data moments to the corresponding model moments. "IQR" refers to the interquartile range (divided by the median) which measures the variability of time use choices.

goods.²³²⁴ Third, the FES data shows that there is also substantial dispersion in the public good share for couples. This empirically supports the notion of heterogeneity in economies of scale, since families who spend a lot on public consumption should have higher economies of scale.

Discussing the implied marriage market patterns is also useful, as these shed some more light on the model mechanisms and estimation. Table 4 describes the extent of assortative matching on preferences based on the estimated model by showing the rank correlations of spouses' preferences. The table shows that there is very little assortative matching on preferences, since the correlations are very small for all three preference dimensions. This is informative about the search incentives for singles in the estimated model: in each period, singles face the choice between starting a relationship with a randomly drawn partner with given characteristics, or rejecting the match, staying single for another period and waiting for another draw from the distribution of partners. The fact that there is little assortative matching on preferences means that the utility gain from matching assortatively on preferences is not as large as the other considerations about potential partners (their wage and the value of the love shock). The extent of assortative matching is important for the analysis, as it determines how much preference heterogeneity there is within couples. Even if there is a lot of heterogeneity in the group of singles (as measured by a high interquartile range in the context of the estimation), this does not automatically mean that there is also a lot of preference heterogeneity between partners, since couples could in theory be matched perfectly assortatively on preferences. The first column of Figure 3 illustrates the joint distribution of couples' preferences graphically by showing the dispersion in preference coefficients.

Note that there is no direct data equivalent for assortative matching on preferences, as preferences are unobserved in the data. The extent of assortative matching is identified from estimating the utility and home production parameters together with the structural assumptions about the marriage market and the parameters of the love shock process. The parameters of the utility and home production function determine the size of the surplus and how it varies with partners' preferences. This in turn determines the gain from searching for a partner with a particular preference type, relative to starting a relationship with the randomly drawn partner in the current period.

²³To make sure that this is not primarily driven by income effects (low income individuals could spend a higher fraction of their expenditure e.g. on housing), I also ran a regression of the share of public consumption on log income, squared log income and year fixed effects. However, this explains only 6.5% of the variance of the expenditure share.

²⁴In the estimation, the variability of public good preferences is estimated from the variation in home production hours. Figure 7 shows one could similarly target the variation in the share of public expenditure, which is also consistent with preference heterogeneity.

TABLE 4: Assortative matching on preferences

	Rank Correlation
Consumption preference (α^C)	0.02
Leisure preference (α^L)	0.03
Home good preference (α^D)	0.04

Notes: The table shows the rank correlations between the preferences of partners in the estimated model.

Another important model implication relates to selection into relationships based on preferences. Theoretically, it could be that individuals with a strong public good preference are more likely to enter relationships in order to benefit from joint home production capacities. It turns out that this selection effect is also very small, which is best illustrated by the second column of Figure 1. For example, panel (b) of the figure shows the distribution of the preference coefficient for the public good (α_i^D), binned into 5 categories, for both coupled individuals and singles. The plot shows that the distributions are very similar, indicating that there is little preference-based selection into relationships. Panel (d) and (f) similarly show the distributions of the leisure and private consumption preference coefficients.

4 Results

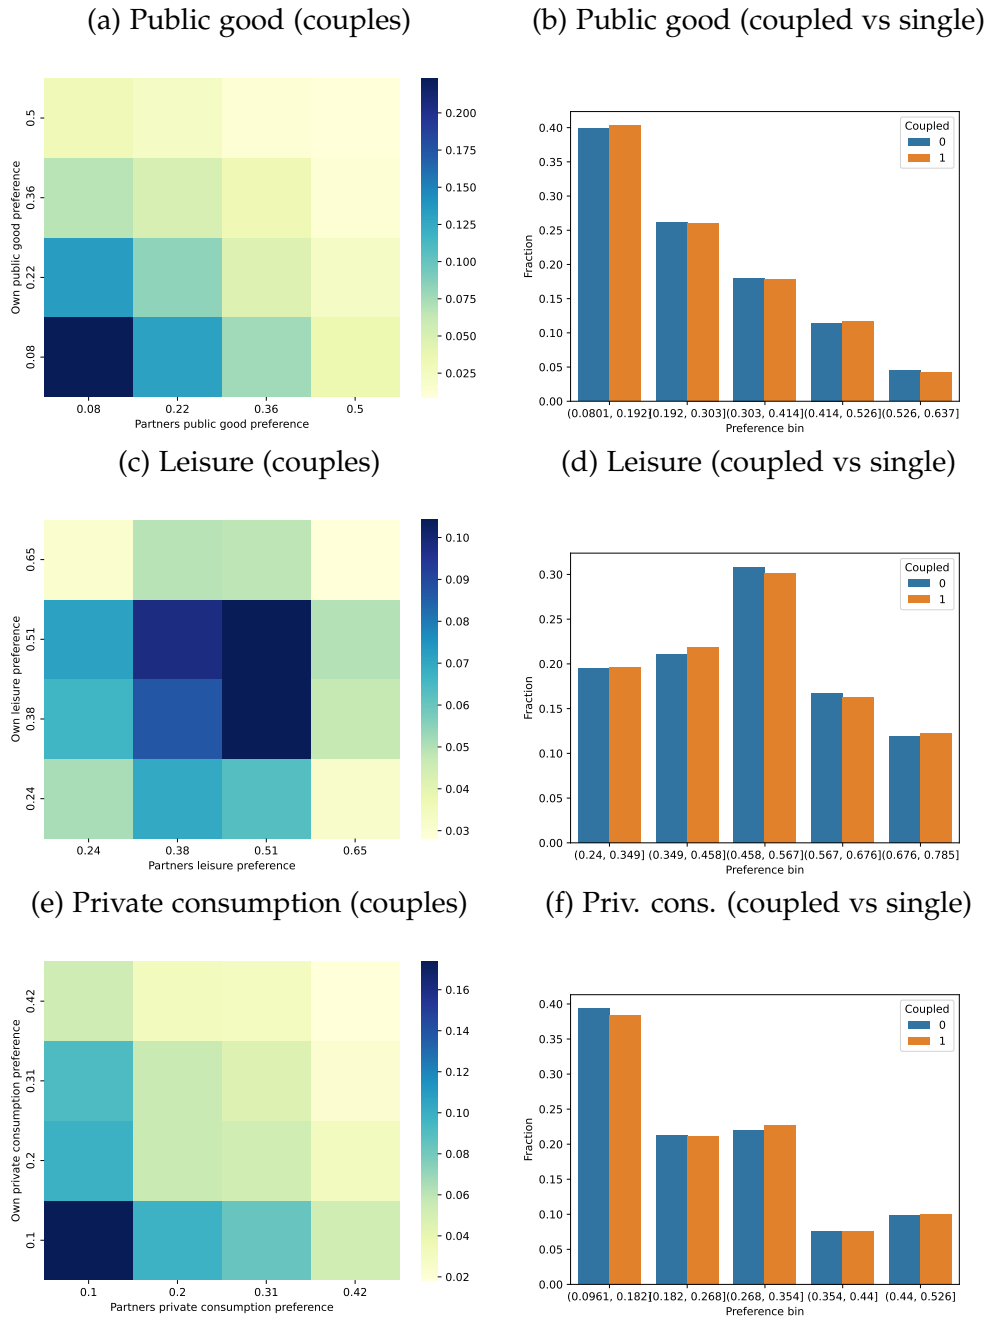
4.1 Welfare Analysis

4.1.1 Household income vs individual welfare

In this section, I turn to the welfare analysis. The main question that the analysis addresses is to what extent household income, i.e. the labour income of the family, is an accurate measure of the individual well-being of each household member, as captured by the *Money-Metric Welfare Index* (MMWI) which was defined earlier. As a starting point for this analysis, I conduct an individual-level poverty analysis, which relates to a large body of work in the collective model literature which has contrasted poverty based on the level of individuals and households (e.g. [Cherchye et al. \(2018\)](#)). To complement this analysis, which focuses on the poorest individuals in the population, I will also discuss overall inequality in the population.

The poverty analysis is conducted using a simulated sample of couples based on the estimated model. An individual is defined as income-poor if their household is in the bottom 20% of the income distribution. Note that income-poverty will always refer to household rather than

FIGURE 1: Preference distributions in the estimated model



Notes: Each panel labeled as *couples* 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.

individual income (that is, the term income-poor will be used for household-income-poor).²⁵ Similarly, *welfare-poverty* is defined as the MMWI of the individual being in the bottom 20% of the distribution of the MMWI. Conceptually, one important point in this analysis is that the quantity of poor individuals is the same for both income and welfare-poverty (20% of the population), but the two measure can have different implications for *who* is considered poor. In other words, the analysis addresses the question to what extent those people who are least well-off in terms of income are also least well-off in terms of the MMWI.²⁶

TABLE 5: Poverty in Household Income vs Welfare

	Not welfare-poor	Welfare-poor
Not HH-income-poor	0.72 (<i>True Negatives</i>)	0.08 (<i>False Negatives</i>)
HH-income poor	0.08 (<i>False Positives</i>)	0.12 (<i>True Positives</i>)

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

This is illustrated by Table 5. Assuming that the ‘true’ poverty concept we are interested in is the MMWI, the table shows the breakdown into true positives, true negatives and classification errors when poverty is analysed based on household income instead. The diagonal shows the fraction of individuals where household income and the MMWI agree on their poverty status: 72% of individuals in the population are not poor according to both measures, and 12% are poor according to both income and welfare. Most importantly, the table shows the fractions of false negatives and false positives: 8% of the population are poor according to household income but not poor according to welfare (false positives) and, similarly, 8% are poor according to welfare but not according to income (false negatives).²⁷ The important implication from this table is that looking at household income does not always indicate who is worst off in terms of welfare. For example, policies which exclusively target the income-poor do not reach all individuals with the lowest levels of welfare and may therefore not be ideally targeted. A useful summary statistic in this regard is the fraction of welfare-poor individuals which is identified when looking at income-poverty. This is the ratio of true positives over all welfare-poor individuals (that is, $0.12/0.20=59\%$). This means that the group of income-poor individuals contains 59% of the welfare-poor or, conversely, that 41% of the welfare-poor are missed when only looking at household income.

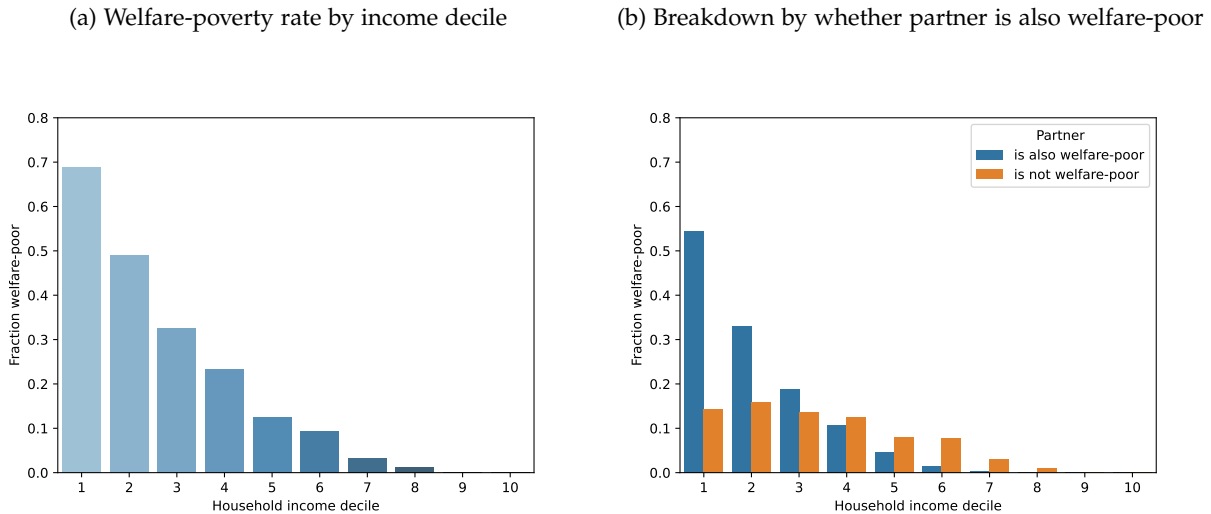
Figure 4 shows the breakdown of welfare-poor individuals by income decile. Panel (a) first

²⁵As families are characterised by at least some degree of sharing, individual income often does not reflect welfare.

²⁶I choose the approach of holding the number of individuals fixed (rather than for example defining a poverty threshold based on income and also applying it to the MMWI) because income and the MMWI are on different scales, since the MMWI is based on full income which also includes the value of time. Therefore, the values of the MMWI are substantially higher than income.

²⁷Note that the number of false negatives and false positives is equal by construction since the number of poor individuals is held constant at 20% for both income and welfare.

FIGURE 2: Poverty in Household Income vs Welfare



Notes: Panel (a) 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. Panel (b) breaks the welfare-poor individuals in each group down into two groups: (1) being welfare-poor *and* having a partner who is also welfare-poor and (2) being welfare-poor but *not* having a welfare-poor partner. Note that the two bars for each income decile in panel (b) add up to the value shown in panel (a) for that decile.

shows the share of individuals in each income decile which is welfare-poor. This figure adds two important points to the analysis. First, it shows more precisely where welfare-poor individuals are located in the income distribution and that a substantial fraction of them is in the middle of the distribution (in the 4th and 5th income decile). Second, the fraction of welfare-poor individuals declines with household income. Looking at the highest income decile, only a very small fraction of people is welfare-poor. This illustrates that income and welfare are positively but imperfectly correlated.

An important part of the analysis will be to investigate the role of (1) *intra-household inequality* and (2) *economies of scale* in more detail. As a first step to illustrate the role of these two factors, panel (b) of Figure 4 breaks down the values from panel (a) into whether the partner of the individual is also welfare-poor. For example, panel (a) shows that 23.2% of individuals in the 4th income decile are welfare-poor. Panel (b) shows that 10.6% are welfare-poor and have a welfare-poor partner and 12.6% are welfare-poor but do not live with a welfare-poor partner. These numbers illustrate that intra-household inequality plays some role: in those households where only one partner is welfare-poor but the other is not, the unequal division of resources within the family contributes to one partner being below the poverty line in terms of welfare. For example, these can be households who have a decent amount of income (since they are in the 4th income decile), but where consumption expenditure is skewed towards one of the partners, leaving the other

with a lower level of welfare. However, panel (b) also shows that there is a substantial number of households who are above the poverty line for income, but where *both* partners are welfare-poor. In these cases, intra-household inequality cannot explain why individuals in medium-income households are welfare-poor. Instead, as will be shown in more detail in section 4.1.3, one factor that plays a role in these cases is heterogeneity in economies of scale. Formally, economies of scale for household h (with members i and j) are defined as $EOS_h = \frac{M_i + M_j}{\bar{w}_i + \bar{w}_j}$ (the ratio of the sum of MMWIs and the full income of the household). Intuitively, this means that some households can easily achieve a high level of welfare for a given level of full income, whereas it is more difficult for others. Due to preference heterogeneity, families where both partners have a high public good preference can spend a lot on the public good and enjoy a high level of welfare, whereas achieving the same level of welfare is more difficult for families where partners have a strong preference for private goods. Looking again at the example of the 4th income decile in panel (b) of Figure 4, 10.6% of individuals are welfare-poor and have a welfare-poor partner, which suggests that low economies of scale play a role in these cases.

4.1.2 The role of consumption, time use and preference heterogeneity

To better understand why welfare, as measured by the MMWI, differs from household income, I perform two decomposition exercises which shed light on different aspects. For the first exercise, it is useful to think about going from household income to individual welfare in three steps. First, one has to take into account that consumption expenditure can be divided unequally within families. Second, in addition to individual consumption, the MMWI also reflects the time allocation of the household. Finally, the MMWI weights consumption and time use using individual preferences. The idea behind the exercise is to account for these additional aspects one-by-one in order to understand the role of these factors.

TABLE 6: What explains the differences between household income and welfare?

	Frac of welfare-poor identified	Rank corr with M_i	Avg rank dev from M_i
Household income	0.59	0.70	17.10
(a) Decomposition 1			
Individual consumption (C_i)	0.56	0.61	19.70
Empirical sharing rule (ESR_i)	0.64	0.83	12.00
MMWI with equal prefs (\tilde{M}_i^{EP})	0.60	0.75	14.88
(b) Decomposition 2			
Full income (\tilde{M}_i^1)	0.74	0.90	9.40
Full income + EOS (\tilde{M}_i^2)	0.80	0.94	6.90
Full income + Sharing (\tilde{M}_i^3)	0.80	0.89	8.60

Notes: This table shows the results from the exercises described in Sections 4.1.2 and 4.1.3. Column 1 shows the fraction of welfare-poor individuals which is identified through a given variable in a row. For example, the first row shows that looking at the income-poor identifies 59% of the welfare-poor. The second row shows the rank correlation between the variable from that row and the MMWI. The third column illustrates the rank correlation more intuitively by showing the average absolute rank deviation between the variable and the MMWI. For the definitions of the variables in each row, see the description in the main text.

For the first step - moving from household income to individual consumption -, we can define can define *total individual consumption* (C_i) as the sum of private (c_i) and public good (q) expenditure:

$$C_i = c_i + q$$

If the difference between household income and the MMWI was primarily driven by consumption inequality within families (but not time use or preference heterogeneity), looking at consumption poverty should identify most of the MMWI-poor individuals. Interestingly, Table 6 shows that looking at consumption does actually *worse* in terms of identifying the welfare-poor; while income identifies 59% of the welfare-poor, consumption only identifies 56%. This shows that intra-household consumption differences alone are not the main factor driving a wedge between income and welfare. Instead, it is also important to look at time use and preferences. In fact, the finding that consumption does worse than income suggests that consumption inequality is driven to a large extent by preference heterogeneity. In this case, low consumption does not necessarily imply a low level of welfare, since the consumption level simply reflects a lower consumption

preference rather than a lower access to household resources.

The second step consists in including time use into the analysis, while still ignoring preferences. 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)). Using an individual's wage rate as a measure of the value of time, the empirical sharing rule takes consumption expenditure ($c_i + q$) and adds the the market value of leisure ($w_i l_i$) and home production time of the individual and their partner ($w_i d_i + w_j d_j$):

$$ESR_i = c_i + q + w_i l_i + w_i d_i + w_j d_j$$

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. Using the empirical sharing rule identifies 64% of welfare-poor individuals. This is an improvement over household income, which identifies 59%, and shows that taking time use into account better characterises welfare, but still the improvement is relatively modest.

Third, and lastly, the role of preference heterogeneity can be assessed by looking at the residual. 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, 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 41% of welfare-poor individuals unidentified.

As an alternative and complementary way of quantifying the role of preference heterogeneity for the analysis, I also recomputed the MMWI under the assumption that each individual has identical preferences, which are given by the sample mean.²⁸ This means computing:

$$\tilde{u}_i(c_i, l_i, D) = \frac{1}{1 - \gamma} \left(\bar{\alpha}^c c_i^\eta + \bar{\alpha}^l l_i^\eta + \bar{\alpha}^D D^\eta \right)^{\frac{1-\gamma}{\eta}}$$

²⁸The reason why this additional exercise is useful is that one drawback of the empirical sharing rule is that it does not take the functional form of utility into account, since it is somewhat unclear what the weighting factors should be when adding up consumption and time use. In the context of the model, the exercise of recomputing the MMWI while assuming equal preferences is the cleanest way of isolating the role of preference heterogeneity. An advantage of the ESR, however, is that it can in principle be computed without the need for a model, and is therefore much easier to take to the data when suitable data is available (such as in the Japanese data used by Lise and Yamada (2019)). Therefore, both exercises are useful.

Here, $\bar{\alpha}^c$, $\bar{\alpha}^l$ and $\bar{\alpha}^D$ are the population means. Then, the MMWI is recomputed as

$$\tilde{M}_i^{EP} = \bar{V}^{-1}\left(\tilde{u}_i(c_i, l_i, D)\right)$$

EP stands for *equal preferences* and \bar{V}^{-1} is the indirect utility function when preferences are given by the population means. The idea behind this exercise is that if there is little dispersion in the preference coefficients, the actual MMWI and \tilde{M}_i^{EP} should be close and \tilde{M}_i^{EP} should be a good predictor of a welfare. In practice, the results from this exercise again emphasise the importance of preference heterogeneity: looking at the poorest 20% in terms of \tilde{M}_i^{EP} only identifies 60% of the welfare poor, which is a slight improvement over household income but leaves an important fraction of the welfare-poor unidentified. Alternatively, this calculation could also be interpreted as estimating a restricted model without preference heterogeneity (where the preferences of every individual are estimated as the population mean) and inferring the MMWI from the restricted model. The result suggests that allowing for preference heterogeneity is important for assessing the MMWI.

Taking stock, the results from this exercise shed light on the relative importance of consumption, time use and preference heterogeneity in driving the difference between the MMWI and household income. The starting point was that income only identifies 59% of welfare-poor individuals. Looking at consumption does slightly worse than income and identifies 56%, whereas looking at the empirical sharing rule (which takes consumption and time use into account) identifies 64%. This finding has two implications. First, these results illustrates that preference heterogeneity, which is an understudied aspect in the literature on individual welfare, is important for the analysis. In other words, the *interaction* between consumption, time use and preferences are key for assessing individual welfare. Second, from a more empirical perspective, it highlights the limitations of drawing welfare conclusions based on data about allocations alone. In principle, the empirical sharing rule is straightforward to directly compute from the data without the need for a model, as long as joint data on consumption and time use is available (such as in the Japanese data used in [Lise and Yamada \(2019\)](#)). However, the results demonstrate that it can be difficult to fully characterise welfare based on such data alone. Therefore, it would be very useful for future data collection efforts to also construct empirical measures of preferences in order to get a better picture of welfare inequality in the population.

4.1.3 The role of intra-household inequality and economies of scale

The second decomposition focuses on the role of intra-household inequality and economies of scale. Another way of thinking about the link between household income and individual welfare is through the lens of the following identity²⁹:

²⁹This can alternatively also be written as $M_i = (y_h + w_i l_i + w_j d_i + w_j l_j + w_j d_j) \cdot \frac{M_h}{\bar{w}_h} \cdot \frac{M_i}{M_h}$.

$$\underbrace{M_i}_{\text{MMWI}} = \underbrace{y_h}_{(0) \text{ HH income}} \cdot \underbrace{\frac{\bar{w}_h}{y_h}}_{(1) \text{ Transform to full income}} \cdot \underbrace{\frac{M_h}{\bar{w}_h}}_{(2) \text{ Add EOS}} \cdot \underbrace{\frac{M_i}{M_h}}_{(3) \text{ Add sharing}} \quad (1)$$

In terms of definitions, note that h is the household index (of individual i and their partner j). $M_h = M_i + M_j$ is the sum of the MMWI of the partners and $\bar{w}_h = \bar{w}_i + \bar{w}_j$ is the full income of the household. On the left-hand side of the equation, we have the MMWI of the individual (M_i). On the right-hand side, we start with household income (y_h) and link it to the MMWI making several adjustments:

1. **Full income:** The first step adjusts for the fact that the MMWI based on full income (taking into account not only household income but also the value of time) by transforming household income into full income (by multiplying income with $\frac{\bar{w}_h}{y_h}$).
2. **Economies of scale:** The next step adjusts for economies of scale. Total well-being of the household can be higher than full income because of economies of scale, which is reflected by the factor $EOS_i = \frac{M_h}{\bar{w}_h}$. Recall that the MMWI of singles is equal to their full income (since the definition of the MMWI is the full income someone who need as a single to achieve a certain living standard). Therefore, if both partners were single, their combined welfare would be $\bar{w}_i + \bar{w}_j = \bar{w}_h$. In the relationship, by contrast, their combined welfare is $M_i + M_j = M_h$. Therefore, $EOS_i = \frac{M_h}{\bar{w}_h}$ is a measure of economies of scale. With $EOS_i > 1$, total welfare in the couple is higher than if partners were single.
3. **Sharing:** Finally, the factor $\frac{M_i}{M_h}$ adjusts for an unequal division of welfare within the household. $\frac{M_i}{M_h} = 0.5$ corresponds to welfare equality within the household; if the fraction deviates from 0.5, one partner is relatively better off.

In words, this identity essentially says that since the MMWI is a concept based on full income, it can be thought of as transforming household income into full income and then additionally applying economies of scale and taking the within-household division of welfare into account.

We can use this identity in order to better understand the role of each factor. The results are also shown in Table 6. As a starting point, recall that looking at household income identifies 57% of the welfare-poor. Next, based on equation 1, I study the role of the differences between income and full income by computing a "hypothetical" MMWI \tilde{M}_i^1 which omits the last two terms of equation 1 (EOS and sharing) and only takes the first two ones into account. This corresponds to the full income of the household:

$$\tilde{M}_i^1 = y_h \cdot \frac{\bar{w}_h}{y_h} = \bar{w}_h$$

Table 6 shows that full income identifies 74% of the welfare-poor. This means that looking at full income is much closer to welfare than looking at income. But full income alone does not fully characterise welfare, since it ignores economies of scale and sharing and still leaves 26% of

the welfare-poor unidentified. Therefore, next, we can look at the role of EOS and sharing. For the role of EOS, I compute another "hypothetical" MMWI based on equation 1 which ignores the sharing factor but includes everything else:

$$\tilde{M}_i^2 = y_h \cdot \frac{\bar{w}_h}{y_h} \cdot \frac{M_h}{\bar{w}_h} = M_h$$

Including EOS does slightly better than full income: now, 80% rather than 74% of the welfare-poor are identified. Finally, the role of intra-household inequality can be assessed by computing:

$$\tilde{M}_i^3 = y_h \cdot \frac{\bar{w}_h}{y_h} \cdot \frac{M_i}{M_h} = \bar{w}_h \cdot \frac{M_i}{M_h}$$

Interestingly, including intra-household inequality on its own is also a very modest improvement; \tilde{M}_i^3 also identifies 80% rather than 74% of the welfare-poor. Taken together, these two results imply that neither including EOS nor intra-household inequality in isolation characterises individual welfare much better than full income. Instead, the interaction between the two factors is important, meaning both how large economies of scale are, and how they are divided between the partners.

4.1.4 Overall inequality

So far, the analysis has focused on poverty, as defined by the least well-off 20% of the distribution. Another important question, from a more macroeconomic angle, is how the issues discussed so far apply to overall inequality. How different is overall welfare inequality from household income inequality? And what is the role of intra-household inequality and economies of scale for overall inequality?

As a first step for this analysis, which is closely connected to the poverty analysis, Table 6 also shows overall rank correlations along the whole distribution. This addresses the question how strongly someone's rank in the income distribution predicts their rank in the welfare distribution. As a baseline, the rank correlation between household income and the MMWI is 0.7. To illustrate the magnitude of this rank correlation more intuitively, the table also shows the average rank deviation between the two measures. Denoting R_i^Y as the income rank of the individual and R_i^M the welfare rank, this is computed as $\frac{1}{N} \sum_{i=1}^N |R_i^Y - R_i^M|$, where $i = 1, \dots, N$ are the individuals in the simulated sample. Household income mispredicts somebody's rank in the welfare distribution on average by 17.1. For example, this implies that when somebody is in the middle of the income distribution, the welfare rank on average deviates by 17.1 from 50. Therefore, the result that household income is a relatively imprecise measure of the MMWI also applies when looking at the distribution more broadly.

The remaining results regarding the rank correlations from Table 6 mirror the main insights from the poverty analysis. Consumption is slightly less correlated with welfare than household income,

whereas the empirical sharing rule and the MMWI under the assumption of equal preferences improve on household income. In the second exercise about economies of scale and intra-household inequality, one result that stands out is that full income is correlated with welfare substantially more strongly than household income; in this case, the correlation goes up to 0.90 from 0.7 and the average rank deviation is only 9.4 instead of 17.1. Including EOS further improves the correlation to 0.94, whereas including the sharing factor marginally decreases it to 0.89.

In addition to looking at the rank correlations, it is also useful to decompose inequality into a within-household and a between-household component to see how much intra-household inequality contributes to overall inequality. There has been limited research on quantifying the contribution of intra-household inequality to overall inequality. Most notably, [Lise and Seitz \(2011\)](#) estimate a collective household model based on British data and find that 25% of consumption inequality in 2000 is due to within-household inequality. Table 7 shows the results for within-between decompositions based on the estimated model. First, regarding consumption inequality, the table shows that a relatively large fraction of inequality in private consumption is due to within-household inequality (35.3%). This fraction goes down to 23.5% for total consumption, reflecting that total consumption also includes the public good expenditure of the household. Therefore, the overall level of intra-household inequality in consumption is close to the estimate of [Lise and Seitz \(2011\)](#). Note that the model has two different sources of intra-household consumption inequality: consumption differences can either be due to bargaining (the Pareto weight) or due to preference heterogeneity. This affects the welfare implications. Consumption inequality which results from bargaining is "bad" from an inequality perspective, since it is the result of the unequal distribution of intra-household decision power. Consumption inequality resulting from preference heterogeneity, by contrast, is "neutral", because the partner who consumes less is not necessarily worse off in terms of welfare if it results from preference heterogeneity. As a simple way of studying the relative importance of these two factors in the model, I also computed the variance decomposition for the subset of couples where both partners have equal preferences. In this case, the share of within-household inequality goes down to 5% for private consumption and 4.3% for total consumption, suggesting that preference heterogeneity accounts for a large fraction of intra-household consumption inequality.

Regarding inequality more broadly, intra-household inequality accounts for 10.6% for the empirical sharing rule and, most interestingly, for 18.1% in the case of the MMWI, which captures overall welfare inequality and for which there is no previous estimate of the fraction of intra-household inequality available. This finding suggests that the contribution of intra-household inequality to overall welfare inequality is economically meaningful and welfare inequality would be understated by looking at between-household inequality only.

The final question is about the role of economies of scale for overall inequality. As discussed earlier, economies of scale (EOS) are heterogenous in the population. Families where both partners have a high public good preferences can pool a lot of resources and have higher EOS than those who place more importance on private goods. Therefore, inequality in economies of scale con-

TABLE 7: Within/Between Decompositions

	Total	Within	Between	Fraction within
Private consumption	0.43	0.15	0.28	35.3%
Total consumption	0.27	0.06	0.20	23.5%
Empirical sharing rule	0.11	0.01	0.10	10.6%
MMWI	0.50	0.09	0.41	18.1%

Notes: This table reports within/between decompositions of the log variance of consumption, the empirical sharing rule and the MMWI.

tributes to total welfare inequality. This effect can also be illustrated by looking at the equation for the second decomposition above and taking logarithms:

$$\log(M_i) = \underbrace{\log(y_h)}_{(0) \text{ HH income}} + \underbrace{\log\left(\frac{\bar{w}_h}{y_h}\right)}_{(1) \text{ Transform to full income}} + \underbrace{\log\left(\frac{M_h}{\bar{w}_h}\right)}_{(2) \text{ Add EOS}} + \underbrace{\log\left(\frac{M_i}{M_h}\right)}_{(3) \text{ Add sharing}}$$

This equation highlights that the variance of $\frac{M_h}{\bar{w}_h}$, which captures heterogeneity in economies of scale, contributes to welfare inequality in terms of the MMWI (M_i). Using the standard variance formula for sums, this gives rise to the following approximation of the variance of welfare (combining the first two terms into full income \bar{w}_h):

$$\underbrace{\text{Var log}(M_i)}_{\text{Welfare inequality}} \approx \underbrace{\text{Var log}(\bar{w}_h)}_{\text{Wage inequality}} + \underbrace{\text{Var log}\left(\frac{M_h}{\bar{w}_h}\right)}_{\text{EOS inequality}} + \underbrace{\text{Var log}\left(\frac{M_i}{M_h}\right)}_{\text{Intra-HH inequality}} + \underbrace{\text{Cov}\left(\bar{w}_h, \frac{M_h}{\bar{w}_h}\right)}_{\text{Cov of full inc and EOS}}$$

This is an approximation because the variance of M_i also depends on the covariances between full income and the sharing term and between EOS and the sharing term.³⁰ The equation says that from the perspective of the MMWI, welfare inequality is closely related to full income (that is, wage) inequality.³¹ However, wage inequality needs to be adjusted for (1) heterogeneity in

³⁰Note that these should be very small since the sample includes both partners for each couple.

³¹Welfare inequality is based on full income/wage inequality because the MMWI is defined in terms of full income. A similar variance decomposition based on household income inequality could be derived with an alternative definition of the MMWI. The choice between these definitions relates to the welfare economics literature (Fleurbaey (2006)). Denote as Y_i an adjusted MMWI, which computes the labour income (rather than the full income) that someone would need to obtain their current utility level. Then (ignoring the covariance term),

$$\underbrace{\text{Var log}(Y_i)}_{\text{Welfare inequality}} \approx \underbrace{\text{Var log}(y_h)}_{\text{Household income inequality}} + \underbrace{\text{Var log}\left(\frac{Y_h}{y_h}\right)}_{\text{EOS inequality}} + \underbrace{\text{Var log}\left(\frac{Y_i}{Y_h}\right)}_{\text{Intra-HH inequality}}$$

M_i and Y_i capture different normative judgements about preference heterogeneity in the spirit of Fleurbaey (2006). Y_i is related to the 'Rente Criterion' from Fleurbaey (2006), which is defined in terms of non-labour (rather than labour)

economies of scale, (2) intra-household inequality and (3) the covariance between full income and EOS. The equation implies that welfare inequality is typically larger than full income/wage inequality, unless high wage households have much lower economies of scale than low wage households. In the context of the estimated model, the covariance between full income and EOS turns out to be very small, so that it does not play a substantial role for the analysis.³²

With this equation in mind, a simple way of quantifying the role of economies of scale for overall inequality is to compute the ratio between the variance of the EOS term and total inequality:

$$\frac{\text{Var} \log\left(\frac{M_h}{\bar{w}_h}\right)}{\text{Var} \log(M_i)} = 23.1\%$$

This means that 23.1% of welfare inequality is due to differences in economies of scale between households. This number is interesting because standard approaches of analysing inequality do not allow for a lot of heterogeneity in economies of scale, as equivalence scales only take the number of household members into account. According to the model, equivalence scales should depend on all characteristics of the household, including preferences and wages, and reflect how much the household spends on public goods.³³ In that sense, equivalence scales are endogenous and depend on household characteristics and optimal choices rather than being a fixed number. This highlights that the assumptions underlying the commonly used equivalence scales can be restrictive, and that structural models can be a useful tool to model equivalence scales more flexibly.³⁴

4.2 Policy Simulation: Minimum Wages and Poverty

To illustrate the usefulness of individual-based welfare measures for policy analysis, I use the estimated model to perform counterfactual experiments relating to minimum wage policy. The minimum wage is an interesting policy in this context because it is often analysed from a household-level rather than an individual-level perspective. A well-known fact is that many minimum wage earners live in households which are above the poverty threshold and are more in the middle of the income distribution (Burkhauser (2015)). Panel (a) of Figure 3 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

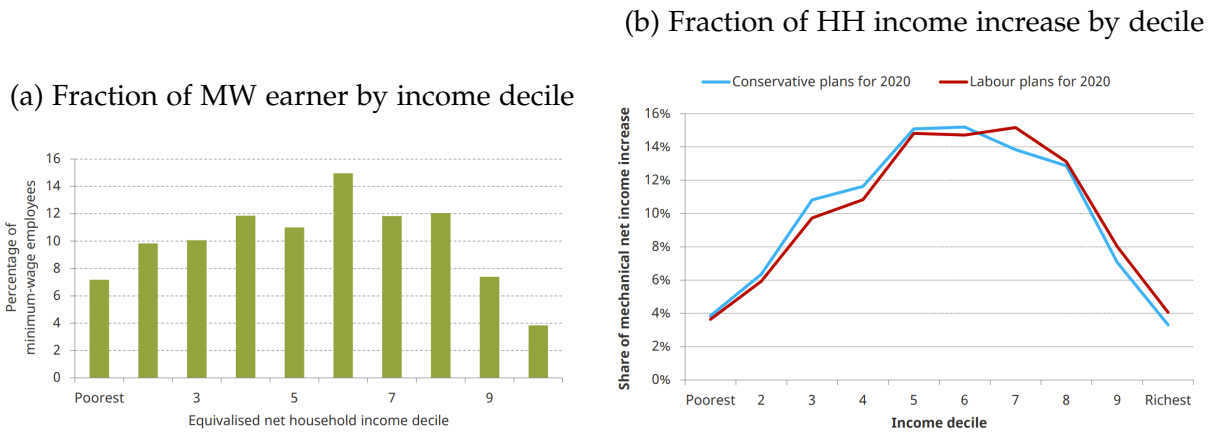
income (see Appendix A).

³²In the model of Salcedo, Schoellman, and Tertilt (2012), this effect would likely play a larger role since they argue that well-off households could have lower economies of scale. In their model, high-income households spend a lower fraction of their budget on public goods than low-income households due to an income effect. As a result, high-income households share fewer resources and have lower economies of scale.

³³As an empirical illustration, Appendix Figure 7 plots the fraction of spending on public goods based on the data from Lise and Seitz (2011) and shows that there is substantial dispersion in public good spending in the data.

³⁴This is related to Salcedo, Schoellman, and Tertilt (2012), who show how the implied equivalence scales from their model vary with income, which is not captured by standard equivalence scales.

FIGURE 3: Minimum Wage Earners along the Income Distribution in the UK



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.

the peak is the 6th decile, with around 14% of minimum wage earners being in this category. Panel (b) of Figure 3 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 households who are not income-poor. From a household perspective, this suggests that the minimum wage may not be an ideally targeted policy to reach the poorest households in the population.

The estimated model allows me to investigate the targeting of minimum wage policy from an individual perspective. In the previous analysis, panel (a) from Figure 4 showed that a substantial fraction of welfare-poor individuals live households which are not income-poor. Therefore, from the perspective of the model, minimum wage earners can be (welfare-)poor even when they do not live in an income-poor household. This makes the question of targeting of minimum wage policy more complicated than suggested by Figure 3, according to which most minimum wage earners are not poor simply by not living in an income-poor household. The goal of the policy experiment is to show that minimum wage increases reduce welfare-poverty among minimum wage earners who live with higher-earning partners, who would not be considered as poor in a conventional income-based sense.

For the experiment, I simulate a 20% increase in the minimum wage. Since the goal of the analysis is to contrast the individual vs household-level effects, rather than to provide a complete ex-ante evaluation of minimum wage policy, I abstract from employment effects.³⁵ Panel (a) of

³⁵In a review of the available evidence, Dube (2019) concludes that such employment effects are likely to be small,

Figure 4 shows the welfare-poverty rate of minimum wage earners depending on the wage of their partner.³⁶ In couples where both partners are MW earners, the welfare-poverty rate is close to 100%, meaning that individuals in these households are certain to be in the bottom quintile of the welfare distribution. The higher the wage of the partner, the lower the probability of being welfare-poor. The part of the distribution which is most important for this analysis is the middle of the distribution, where a minimum earner is in a relationship with a partner who earns a medium wage, such as around £20-25. In these cases, still around 50% of the minimum wage earners are welfare-poor in the status quo. In the experiment, I hold the threshold for welfare-poverty constant at its value from the status quo (where it is defined such that 20% of the population are welfare-poor). The minimum wage increase reduces poverty conditional on this value of the threshold. Panel (a) shows that at the very bottom of the distribution, poverty rates are unchanged and still at 100%, as the MW increase does not push these households over the threshold. The decrease in poverty is most pronounced in the middle of the distribution, where partners' wage is between £20-25. Connecting this back to Figure 3, these households are likely to be in the middle of income distribution, given that the partner earns a fairly good wage in these cases. Overall, the minimum increase reduces welfare-poverty of minimum wage earners, particularly in unequal households where a minimum wage earner is in a relationship with a higher earner.

To conclude the analysis, I discuss the underlying mechanisms for change in welfare-poverty among minimum wage earners in more detail. Following the minimum wage increase, there are three effects which affect welfare poverty rates. The first effect is the *direct effect*. Suppose that the distribution of couples and labour supply choices (and other time use choices) are held constant when the minimum wage increases. The higher minimum wage increases the disposable income of households and each couple can spend more on partners' personal consumption and the public good. To allocate the additional income across both partners and the public good (without allowing the household to reoptimise their choices), I compute the fraction of total expenditure of each good under the status quo and distribute the additional resources according to these shares.³⁷ Essentially, the direct effect increases welfare since it increases income for all households which include MW earners, and therefore these households will spend more on consumption goods. This is shown in panel (b) of the figure, where the blue (solid) line shows the change in welfare-poverty induced by the direct effect.

However, the direct effect does not capture the entire effect of minimum wage policy. After the minimum wage increase, relative wages within couples have changed and there are also income

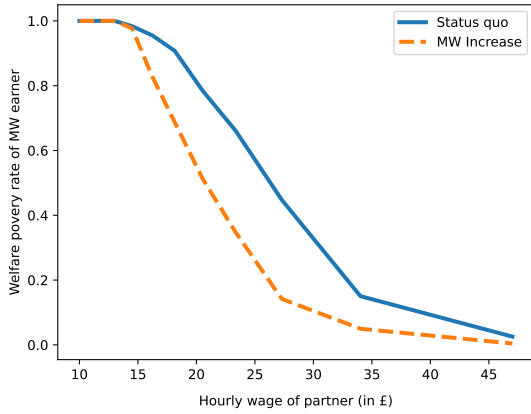
at least for moderate minimum wage increases. It would be possible to include supply-side effects in this model, for example by having a type-specific frictionless labour market with market clearing, but this would make the model more complicated.

³⁶Note that the x-axis shows the wage of the partner, rather than household income like in Figure 3. Since wages are exogenous, it is easier to compare individuals based on wages between the two steady states than based on income (where couples have different positions in the income distribution before and after the policy change).

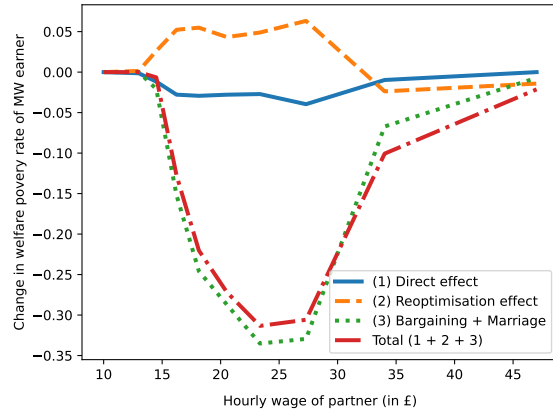
³⁷For example, in a couple in which 30% is spent on partner A, 40% on partner B and 30% on the public good, the extra budget is spent accordingly.

FIGURE 4: Minimum Wage (MW) Analysis


(a) Welfare-poverty rate of MW earners by partners' wage



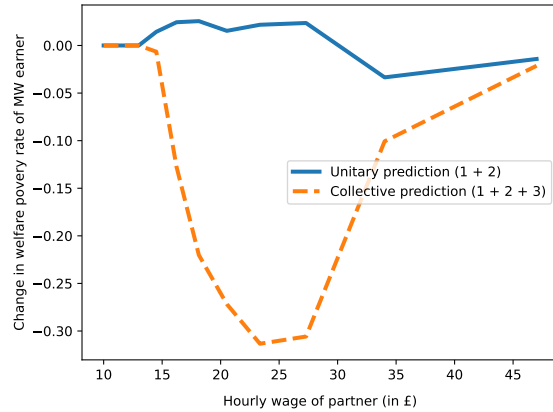
(b) Decomposition of mechanisms



(c) Change in Pareto Weights



(d) Unitary vs collective prediction



Notes: Panel (a) shows the fraction of minimum wage earners who are welfare-poor (i.e. in the lowest 20% of the distribution of the MMWI) depending on the wage of their partner, both for the status quo and the policy experiment where the minimum wage is increased by 20%. Panel (b) shows the decomposition of mechanisms, where the total change is decomposed into the direct effect, reoptimisation effect and bargaining and marriage effect (see main text for the definition of these effects). Note that the total change in panel (b) corresponds to the difference between the lines in panel (a). Panel (c) shows the change in Pareto weights between the status quo and the policy experiment. Panel (d) shows the decomposition from panel (b), but grouped differently: the 'unitarity' prediction is the sum of the direct and reoptimisation effect (channel 1 + 2) and the 'collective' prediction is the sum of all three channels.

effects. Couples will reoptimise their choices (that is, decide on time use and consumption according to the household decision problem). The second effect is therefore the *reoptimisation* effect. I compute this effect by still holding the distribution of couples constant as under the status quo, but solving the decision problem of each household again given the new budget constraint. Panel (b) of the figure plots the reoptimisation effect, expressed as the difference between the poverty rate after reoptimisation relative to the poverty rate after accounting for the direct effect.³⁸ Interestingly, the reoptimisation effect *increases* welfare-poverty among minimum earners when they are in a relationship with a higher-earning partner. The reason for this effect is that because the minimum wage earner now has a higher relative wage, the household will allocate more work hours to them. As a result, leisure and therefore welfare go down.

Finally, the third step is the *bargaining and marriage effect*. So far, the distribution of couples and the Pareto weights were held constant. As the minimum wage increase shifts the relative wage within the couple more in favour of the minimum wage earner, this will increase their intra-household decision power (the Pareto weight). This effect is shown in panel (c) of the figure. In the status quo, the Pareto weight of the minimum wage earner declines when the wage of the partner increases and the increase in the minimum increases the Pareto weight of the minimum wage earner in unequal couples. This effect increases their consumption and it also gives them more leisure (at least partially counteracting the loss of leisure from the reoptimisation effect). In addition, the composition of couples can also differ between the two steady states, which is also included in this step.³⁹ Panel (b) of the figure shows that the bargaining and marriage effect plays a substantial role in generating the reduction in welfare-poverty among minimum wage earners, as it offsets the other two effects and reduces welfare-poverty by increasing the Pareto weight of the minimum wage earner, suggesting that the intra-household reallocation of welfare is one of the key mechanisms in unequal couples.

To summarise this discussion, the minimum wage increase has three effects. First, the direct effect increases welfare of minimum wage earners because they now earn more and can spend more on consumption. Second, the reoptimisation slightly offsets the direct effect. Since minimum wage earners lose leisure due to changes in relative wages, their welfare-poverty rates increase. Finally, the bargaining and marriage effects leads to a substantial reduction in welfare-poverty rates when minimum wage earners are in a relationship with a higher earner.

One interesting take-away message from this analysis relates to the comparison between the implications of 'unitary' and 'collective' models of the household. The unitary model has often been used in the literature to describe household behaviour (see e.g. [Chiappori and Mazzocco](#)

³⁸More precisely, denote as p_j^0 the poverty rate of minimum wage earners with partner wage j in the status quo. p_j^1 is the poverty rate when allowing for the direct effect. p_j^2 is the poverty rate after reoptimisation given the new budget constraint and p_j^3 is the poverty rate in the new steady state with the adjusted minimum wage policy. The direct effect is $p_j^1 - p_j^0$, the reoptimisation effect is $p_j^3 - p_j^2$ and the bargaining and marriage effect is $p_j^3 - p_j^2$.

³⁹Since bargaining and marriage are determined together it is less straightforward to separate the effect of changes in the distribution of couples and changes in intra-household bargaining power.

(2017) for a survey of the literature). Chiappori and Mazzocco (2017) define unitary models as having a stable utility functions which does not depend on prices, income and distribution factors such as relative wages, irrespective of the number of household members. Recall that the household decision problem from the model consists in maximising a weighted sum of the individual utility functions of the partners:

$$\lambda_{ht}u_i(c_i, l_i, D) + (1 - \lambda_{ht})u_j(c_j, l_j, D)$$

In a unitary version of the model, the Pareto weight λ_{ht} would be exogenously fixed (for example to $\lambda_{ht} = 0.5$), which would result in a stable utility function. In a collective model, the Pareto weight is an endogenous parameter, which for example reflects relative wages of the partners. Since microeconomic studies typically reject the unitary model in favour of collective models, an important topic is to better understand the differences between these two types of frameworks for policy analysis.

Looking at panel (b) from Figure 4, the figure shows that the implications between the two types of models are very different regarding the impact of minimum wage policy: a unitary model would only capture the direct and reoptimisation effect, but not the bargaining effect. As a result, the prediction from a unitary model would be the opposite of what happens in the full model: since the reoptimisation effect is stronger than the direct effect, raising minimum wages would overall lead to a small *increase* in welfare-poverty among minimum wage earners. This is shown as the blue (solid) line in panel (d), which sums the direct and reoptimisation effect from panel (b). As discussed above, the key mechanism here is that conditional on a Pareto weight, the higher wage of the minimum wage earner increases the opportunity cost of their leisure, so that the household reallocates work hours towards them. By contrast, allowing for Pareto weights to adjust overturns this effect and leads to a *reduction* in welfare-poverty among minimum wage earners when they are in a relationship with a higher earner (the orange/dashed line in panel (d)). Therefore, unitary and collective models have different implications for the analysis of minimum wage policy, which should be kept in mind when modelling the impact of minimum wages on family labour supply. Given the microeconomic evidence in favour of collective models, a unitary model could lead to misleading conclusions.⁴⁰

⁴⁰This is related to the point made by Knowles (2012), who shows that the decline of the gender wage gap (which increases women's relative wage) has very different implications in a unitary and a collective model and that ignoring the collective effect has strong implications for the model's prediction about relative leisure.

5 Conclusion

In this paper, I have studied the link between household income and individual welfare from the perspective of a rich collective household framework which jointly accounts for consumption inequality within households, heterogeneity in time use and individual-level preference heterogeneity. I estimate the model using British data and use it for an individual-level welfare analysis, where I compare poverty and inequality based on household income and on the *Money-Metric Welfare Index* (MMWI), which was proposed by [Chiappori and Meghir \(2015\)](#).

The first main result is that that poverty in household income is an imperfect predictor of individual-level welfare poverty. Welfare poverty is based on considering an individual's total utility from consumption, leisure and the home good. Income poverty only identifies 59% of the welfare-poor, meaning that a substantial fraction of welfare-poor individuals is missed in that case. I also find that welfare-poverty is difficult to predict without taking information on preferences into account, using either only information on consumption or combining consumption and the value of time into a single index. This highlights that unobserved preference heterogeneity, an understudied factor in the literature on individual welfare analysis, is a key component of the analysis. Regarding overall inequality, intra-household inequality accounts for 18.1% in a variance decomposition of individual, and heterogeneity in economies of scale account for 23.1%. Finally, to illustrate the role of individual-level welfare measures for policy analysis, I perform a policy experiment where I increase the minimum wage and study how this affects welfare-poverty among minimum wage earners depending on the position of their partner in the wage distribution. This policy change reduces welfare-poverty particularly for minimum wage earners who have a higher-earning partner who is in the middle of the wage distribution.

The key implication from these findings is that the differences between household income and the *Money-Metric Welfare Index* (MMWI) can be quantitatively important when analysing poverty and inequality. In particular, the *interaction* between consumption sharing, time use and preference heterogeneity is critical for assessing individual welfare. These results point into several directions for future research aimed at improving measurement of the MMWI. From an empirical perspective, my results suggest that it would be important to collect information on individual consumption, time use and individual preferences in a joint dataset, since all of these aspects are needed to draw conclusions about welfare. At present, datasets which jointly measure consumption and time use are still rare, and no direct data is available when it comes to the preference parameters needed in collective models. An important question in this regard is how these preferences can best be measured, as it is less clear whether for example leisure preferences 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.

On the theoretical side, there 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.

References

- Alon, Titan, Matthias Doepke, and Sena Ekin Coskun. 2018. "Trends in Work and Leisure: It's a Family Affair." In *2018 Meeting Papers*, 552. Society for Economic Dynamics.
- Atkinson, Anthony B and François Bourguignon. 2014. *Handbook of income distribution*, vol. 2. Elsevier.
- Attanasio, Orazio P and Luigi Pistaferri. 2016. "Consumption inequality." *Journal of Economic Perspectives* 30 (2):3–28.
- Balleer, Almut, Monika Merz, and Tamás K Papp. 2021. "Couples' Time-Use and Aggregate Labor Market Outcomes." .
- 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.
- Boerma, Job and Loukas Karabarbounis. 2020. "Labor market trends and the changing value of time." *Journal of Economic Dynamics and Control* 115:103885.
- 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 under-nutrition 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 Demuyne, 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-André 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.
- Dube, Arindrajit. 2019. "Impacts of minimum wages: review of the international evidence." *University of Massachusetts, Amherst. <https://www.gov.uk/government/publications/review-of-the-international-evidence-on-the-impacts-of-minimum-wages/review-of-the-international-evidence-on-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.
- Fang, Lei, Anne Hannusch, and Pedro Silos. 2021. "Luxuries, necessities, and the allocation of time." .
- 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.
- Gayle, George-Levi and Andrew Shephard. 2019. "Optimal taxation, marriage, home production, and family labor supply." *Econometrica* 87 (1):291–326.
- Goussé, Marion, Nicolas Jacquemet, and Jean-Marc Robin. 2017. "Marriage, labor supply, and home production." *Econometrica* 85 (6):1873–1919.
- 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. 2021. "Marriage, Labor Supply and the Social Safety Net." *Unpublished manuscript* .
- Mazzocco, Maurizio, Claudia Ruiz, and Shintaro Yamaguchi. 2013. "Labor supply, wealth dynamics, and marriage decisions." *Unpublished manuscript* .

- Obermeier, Tim. 2019. "The Marriage Market, Inequality 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." .
- Salcedo, Alejandrina, Todd Schoellman, and Michèle Tertilt. 2012. "Families as roommates: Changes in US household size from 1850 to 2000." *Quantitative Economics* 3 (1):133–175.
- Shephard, Andrew. 2019. "Marriage market dynamics, gender, and the age gap." *Unpublished manuscript* .
- Shimer, Robert and Lones Smith. 2000. "Assortative matching and search." *Econometrica* 68 (2):343–369.
- Tommasi, Denni. 2019. "Control of resources, bargaining power and the demand of food: evidence from PROGRESA." *Journal of Economic Behavior & Organization* 161:265–286.

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.^{41,42} The goal of this section is to demonstrate that the MMWI can also be viewed through the framework of Fleurbaey (2006) and to discuss 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 individual's 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 the solution to the following hypothetical optimisation problem:

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

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?"

⁴¹Also see e.g. Fleurbaey and Maniquet (2018), Fleurbaey et al. (2008) or Fleurbaey and Maniquet (2011).

⁴²These 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.⁴³ From this perspective, the MMWI is part of a broader family of welfare measures which take preference heterogeneity into account in different ways.⁴⁴ 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 different notions of justice. These are best illustrated by looking at an example in the context of singles. Table 8 shows the choices and welfare criteria for three hypothetical singles who differ in their preference vector p_i :

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

The first single is "consumption-oriented", meaning that they place a larger importance on private consumption than on leisure or home production. Similarly, the second and third individual are leisure- and 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 off 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

⁴³Chiappori 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.

⁴⁴Fleurbaey (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.

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.

TABLE 8: Welfare Measures - Examples (Singles)

Preferred Good	Consumption (<i>c</i>)	Leisure (<i>l</i>)	Home (<i>D</i>)
(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.

To what extent does the choice between the MMWI and the Rente Criterion affect the main results from the analysis, i.e. the welfare analysis of couples? Figure 5 first shows the breakdown of welfare-poor individuals by household income decile, which was one of the main results from the analysis. The figure shows that both the MMWI and the Rente Criterion lead to a fairly similar picture in this regard. For both welfare measures, there is a substantial fraction of the welfare-poor who do not live in income-poor households. According to the Rente Criterion, the fraction of welfare-poor individuals is slightly larger in the lowest income decile.

While the overall rates of welfare-poverty are fairly similar for the two criteria, there is more disagreement about *who* is classified as welfare-poor. Table 9 shows how welfare-poverty accord-

ing to each criterion correlates with the exogenous characteristics of individuals (preferences and wages). Labour market ability (that is, the wage) is negatively correlated with poverty for both criteria, although somewhat more strongly for the MMWI. In terms of preferences, the correlations differ more strongly between the MMWI and the Rente Criterion: Rente Criterion poverty is much more strongly negatively correlated with a low consumption preferences, and positively correlated with the leisure preference of the individual. 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.


Due to these differences, looking at poverty according to the Rente Criterion identifies only 71% of individuals who are poor according to the MMWI (and vice versa, as shown in Table 10). This number illustrates that the conclusion about who is welfare-poor somewhat depends on the welfare measure and the ethical judgements associated with it. The table also replicates the first decomposition exercise from the main text, which goes from income to welfare in three steps (consumption, time use and preferences). Recall that the main result for the MMWI was that looking at income identifies 59% of welfare-poor individuals, looking at consumption does slightly worse and identifies 56%, and the empirical sharing rule (which includes the value of time) does better and identifies 64%. For the Rente Criterion, these results are different: consumption-poverty is the *better* proxy for welfare-poverty and the empirical sharing rule performs *worse*. This is driven by the correlation between preferences and the Rente Criterion, since individuals with a low consumption preference are more likely to be poor according to the Rente Criterion. 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.

TABLE 9: Correlation between poverty, preferences and wages

	MMWI	Rente Criterion
Characteristic		
Consumption pref. (α_i^C)	-0.15	-0.42
Leisure pref. (α_i^L)	-0.06	0.13
Home Good pref. (α_i^D)	0.18	0.22
Ability (a_i)	-0.41	-0.29

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

FIGURE 5: 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 10: Relationship between all welfare measures

True Welfare Measure	Frac of poor identified	
	MMWI	Rente Criterion
Chosen Welfare Measure		
(1) Income	0.59	0.64
(2) Consumption	0.56	0.70
(3) Empirical sharing rule	0.64	0.45
(4) MMWI	1.00	0.71
(5) Rente Criterion	0.71	1.00

Notes: This table shows the relationship between the poverty according to the different variables. The first column shows what percentage of the MMWI-poor is identified by each of the variables in the rows. The second column shows what fraction of the people who are poor according to the Rente Criterion are identified.

Appendix B: Computational Details

Overview of the solution algorithm

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 these state variables (using a grid for the Pareto weight λ and BFGS to find the optimal choices):

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

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 variable 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 ($\Lambda_{t,g}^0(\omega^S)$).
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: $\Lambda_{t,g}^1(\omega^S) = \alpha \Lambda_{t,g}^*(\omega^S) + (1 - \alpha) \Lambda_{t,g}^0(\omega^S)$
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.


Appendix C: Additional Tables and Figures

FIGURE 6: Illustration of preference distribution



Notes: This figure illustrates the preference distribution. Individual preferences consist of a vector $(\alpha_i^C, \alpha_i^L, \alpha_i^D)$ and the three preference coefficients need to add up to 1. This effectively gives rise to a two-dimensional distribution on a triangular space, as shown here.

FIGURE 7: Dispersion in fraction of expenditure spent for public goods



Notes: This figure shows the dispersion of the fraction of expenditure spent for public goods based on the data from [Lise and Seitz \(2011\)](#) (Family Expenditure Survey, 1968-2001). Public expenditure is defined as housing, heating and household durables. Panel (a) shows the fraction of public expenditure for childless singles and panel (b) for childless couples.

CENTRE FOR ECONOMIC PERFORMANCE
Recent Discussion Papers

1953	Kirill Borusyak Xavier Jaravel	The distributional effects of trade: Theory and evidence from the United States
1952	Ariela Caglio Sebastien Laffitte Donato Masciandaro Gianmarco Ottaviano	Has financial fair play changed European football?
1951	Eugenie Dugoua Todd D. Gerarden	Induced innovation, inventors and the energy transition
1950	Christian Krekel Johannes Rode Alexander Roth	Do wind turbines have adverse health impacts?
1949	Réka Juhász Claudia Steinwender	Industrial policy and the great divergence
1948	Virginia Minni	Making the invisible hand visible: Managers and the allocation of workers to jobs
1947	Eugenie Dugoua	Induced innovation and international environmental agreements: Evidence from the ozone regime
1946	Marco Bertoni Gabriel Heller-Sahlgren Olmo Silva	Free to improve? The impact of free school attendance in England
1945	Matthias Mertens Bernardo Mottironi	Do larger firms exert more market power? Markups and markdowns along the size of distribution
1944	Ria Ivandić Anne Sophie Lassen	Gender gaps from labor market shocks
1943	Ihsaan Bassier Alan Manning Barbara Petrongolo	Vacancy duration and wages

1942	Cheng Keat Tang Stephen Gibbons	Are friends electric? Valuing the social costs of power lines using house prices
1941	Luis Bauluz Pawel Bukowski Mark Fransham Annie Seong Lee Margarita López Forero Filip Novokmet Sébastien Breau Neil Lee Clément Malgouyres Moritz Schularick Gregory Verdugo	Spatial wage inequality in North America and Western Europe: Changes between and within local labour markets 1975-2019
1940	Felipe Carozzi Edward Pinchbeck Luca Repetto	Scars of war: The legacy of WW1 deaths on civic capital and combat motivation
1939	Christos Genakos Costas Roumanias Tommaso Valletti	Is having an expert “friend” enough? An analysis of consumer switching behavior in mobile telephony
1938	Andrew E. Clark Maria Cotofan	Are the upwardly mobile more left-wing?
1937	Priya Manwaring Tanner Regan	Public disclosure and tax compliance: Evidence from Uganda
1936	Philippe Aghion Celine Antonin Luc Paluskiewicz David Stromberg Raphael Wargon Karolina Westin	Does Chinese research hinge on US co-authors? Evidence from the China initiative
1935	Nicholas Bloom Steven J. Davis Stephen Hansen Peter Lambert Raffaella Sadun Bledi Taska	Remote work across jobs, companies and space