

Discussion Paper Series – CRC TR 224

Discussion Paper No. 109
Project A 03

The Marriage Market, Inequality and the Progressivity of the
Income Tax

Tim Obermeier *

July 2019

* University of Mannheim, Department of Economics, L7 3-5, 68161 Mannheim, Germany
(Email: tim.obermeier@gess.uni-mannheim.de).

Funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation)
through CRC TR 224 is gratefully acknowledged.

The Marriage Market, Inequality and the Progressivity of the Income Tax*

Tim Obermeier
University of Mannheim

July 17, 2019

Abstract

This paper studies how the progressivity of the income tax affects intra-household inequality and the marriage market. Tax progressivity increases the after-tax earnings of the lower-earning spouse and improves their bargaining position in marriage. This mechanism reduces inequality in consumption and leisure *within* households. In addition, tax progressivity can change who is single and who marries whom. I study these effects in an equilibrium search and matching model with intra-household bargaining, labor supply and savings. The model is calibrated to data from the Netherlands and used to study a hypothetical reform which increases progressivity by 40% relative to the current system. The reduction of intra-household inequality accounts for 24.77% of the reduction in inequality in private consumption due to the reform, and 11.43% of the reduction in inequality in utility from private and public consumption, leisure and home production. Changes in the composition of couples and singles, due to endogenous marriage and divorce, have small implications for inequality.

Keywords: Tax Progressivity, Intra-Household Bargaining, Equilibrium search, Assortative matching

*I'm very grateful to Eckhard Janeba and Michèle Tertilt for their support and invaluable advice. I further want to thank Richard Blundell, Monica Costa Dias, Hanno Foerster, Eric French, Hans-Martin von Gaudecker, Andreas Gulyas, Anne Hannusch, Valérie Lechene, Fabien Postel-Vinay, Krzysztof Pytka, Konrad Stahl, John Wilson, Minchul Yum and many participants at the Macro Lunch and the Public Finance Seminar at the University of Mannheim, the IIPF annual congress 2018 in Tampere and the 2nd CRC TR 224 conference in Montabaur for their helpful comments, as well as the bwHPC team for technical support. Funding by the German Research Foundation (DFG) through CRC TR 224 (Project A03) is gratefully acknowledged. University of Mannheim, Department of Economics, L7 3-5, 68161 Mannheim, Germany (Email: tim.obermeier@gess.uni-mannheim.de).

1 Introduction

A progressive income tax is a common policy tool to reduce inequality in consumption and welfare. The debate on tax progressivity has largely focused on redistribution *across* families and has paid little attention to the question how the allocation of consumption and time *within* families could be affected.¹ Recent empirical work finds that there is substantial inequality in consumption and leisure within households and that intra-household allocations depend on the relative income of spouses (Lise and Seitz (2011), Lise and Yamada (2018)). This suggests that in marriages where one spouse has a higher earnings potential than the other, increasing tax rates at higher incomes and lowering them for lower incomes could improve the bargaining position of the lower-earnings spouse. This bargaining effect would reduce intra-household inequality in consumption and leisure, on top of the usual effect of redistributing from richer to poorer households. In addition, progressivity could affect inequality through its impact on marriage and divorce. Whether an individual is single or married, and whom they are married to, has a significant impact on their living standard. How does tax progressivity affect intra-household inequality and the marriage market? And how much do these channels affect inequality relative to the usual effects of progressivity?

To address these questions, I develop an equilibrium search and matching model with intra-household bargaining, labor supply and savings. Wage inequality results from differences in labor market ability and shocks over time. Couples benefit from joint consumption, can pool home production time and share labor market risk, but not all goods are shared equally between spouses. Personal consumption and leisure are private goods. Couples decide efficiently on allocations, but are unable to fully commit to future allocations (as in e.g. Mazzocco, Ruiz, and Yamaguchi (2013)). The bargaining position of a spouse is given by how well-off they would be as single. Bargaining positions are linked to the marriage market, since being single includes the possibility of marrying in the future. In the example of a high-wage husband married to a low-wage wife, the husband would be able to afford a higher living standard if he was single and would have better future marriage prospects. In this case, a progressive income tax improves the outside option of the wife, since it increases the after-tax income she would earn as single. This raises her share of consumption and leisure in marriage and reduces intra-household inequality.

Tax progressivity further affects marriage and divorce, besides the sharing of resources in marriage. When meeting on the marriage market, individuals have to decide whether to get married or to stay single and wait to meet someone else. They take both the economic characteristics of their potential partner and a non-economic match quality into account. The main mechanisms through which tax progressivity affects marriage and divorce are that it influences the economic value of individuals on the marriage market and how selective they are about potential partners. An increase in progressivity makes low-income individuals richer and they become more attractive on the marriage market and can afford to be more selective. The effects on high income

¹See e.g. Heathcote, Storesletten, and Violante (2017) or Conesa and Krueger (2006).

individuals are the opposite. In addition, from the point of view of each individual, progressivity makes all potential partners more similar in their economic characteristics and reduces the value of staying single in order to continue search.

I calibrate the model to Dutch survey data on intra-household allocations, labor market outcomes and marital histories. Before turning to the policy analysis, I first decompose cross-sectional inequality in private consumption and welfare for the status quo. I decompose the population variance into contributions due to inequality within and between married couples, and due to singles. The within-couple component captures that spouses consume different amounts of private goods, which is typically absent in studies of progressivity. In the calibrated model, the within-component accounts for 23.9% of the total variance of private consumption, singles for 21.4% and the between-couple component for 54.3%. The relative magnitude of the within- and between variance in the model compares well to the data on private consumption from the Dutch LISS panel. On top of private consumption, the calibrated model allows to study inequality in the *utility* from private and public consumption, leisure and home production. The advantage of this measure is that it allows to aggregate welfare from consumption expenditure and time use and takes all public and private goods that are available into account. The contribution of the within-couple variance to total inequality is 6.1% in this case. This reflects households spending part of their time and expenditure on public goods, which are consumed equally by both spouses. As a result, the contribution of the within-couple variance is smaller than for private consumption only.

I then study the effects of a reform that increases in progressivity and analyze the impact on inequality within and across households. The main experiment is a budget-balanced increase in progressivity, by 40%, which increase average tax rates at higher incomes and decreases them at lower incomes.² As the majority of OECD countries, the Netherlands has a system of individual taxation, in which spouses are taxed based on their individual earnings.³ The focus of the analysis is to decompose the total reduction in inequality in terms of changes in within- and across-household inequality. The reform decreases the variance of private consumption by 10.9% and reduces both the within- and the between- household variance. The contribution of the within-household component to the total reduction in inequality is 24.77% for private consumption and 11.93% for the utility measure. Thus, in terms of utility from private and public goods, 11.93% of the total reduction in inequality would be ignored by abstracting from intra-household inequality.

The increase in progressivity in principle affects inequality both through its effects on bargaining positions and on marriage rates and assortative mating. Since there are economies of scale in consumption and home production, married individuals have a higher living standard than singles. Thus, changes in marriage and divorce contribute to the effect of the reform on inequality. In particular, inequality across couples could be affected by a change in the composition of couples. I

²The change in progressivity refers to the progressivity parameter of the tax function used in [Heathcote, Storesletten, and Violante \(2017\)](#) or [Holter, Krueger, and Stepanchuk \(2019\)](#). For details, see section 3.3.

³Exceptions are Germany, the United States and France where tax liabilities are computed based on the income of both spouses.

further use the model to quantify the role of these composition effects. The model rationalizes the observed gender imbalance in domestic work hours with a gender asymmetry in the home production technology, which makes women more likely to reduce their market hours in marriage. This is reflected in matching patterns and low-wage men have a lower probability of getting married, relative to higher wage men. The increase in tax progressivity induces a small increase in the probability than low-wage men get married to higher-wage women. This mechanism alone would generate an decrease in assortative mating. I further decompose the variance across couples to investigate whether such composition effects have a noticeable impact on inequality and find that their role is very small. Composition effects explain between -0.3% and -2.8% of the reduction in inequality within and across couples.

In addition to cross-sectional inequality, I further study how the reform affects expected lifetime welfare. I isolate the role of the marriage market channels (marriage and bargaining) by also computing the welfare effect of the reform for an alternative version of the model, in which marriage and divorce decisions and the intra-household bargaining weight are fixed to the pre-reform decision rules. This can be interpreted as the impact of the reform if marriage market decisions were exogenous and would not react to the reform. I then compare the welfare effect of the policy between the full and the alternative model. The contribution of the marriage market is the difference between these two cases. Without the marriage market adjustments, the reform leads to a welfare reduction by 1.22% , in terms of consumption equivalents relative to the status quo, whereas welfare decreases by only 0.9% when taking the endogenous marriage market responses into account.

The literature on tax progressivity has mostly focused on inequality across households and abstracted from intra-household inequality and endogenous marriage and divorce. For example, in the macroeconomic literature, [Conesa and Krueger \(2006\)](#) or [Heathcote, Storesletten, and Violante \(2017\)](#) quantitatively characterize the optimal level of tax progressivity. Similarly, a large literature in public economics (recently surveyed by [Piketty and Saez \(2013\)](#)) has studied optimal non-linear tax schedule in models without intra-household inequality. A few papers have studied the effects of tax policy in two-earner models (e.g. [Guner, Kaygusuz, and Ventura \(2011, 2012\)](#), [Bick and Fuchs-Schündeln \(2017\)](#)), but typically take intra-household allocations and marriage outcomes as given. Notable exceptions, where taxes are allowed to affect intra-household allocations, are [Alesina, Ichino, and Karabarbounis \(2011\)](#) and [Bastani \(2013\)](#), who study gender-based taxation and thus a different policy than tax progressivity. [Gayle and Shephard \(2019\)](#) study the optimal taxation of couples in a static matching model, in which taxes can affect intra-household allocations, and focus on the optimal degree of jointness of the tax schedule. My paper contributes by focusing on progressivity and the relative importance of inequality within and across households. Thus, my paper sheds additional light on the role of intra-household inequality. In addition, I analyze a dynamic marriage market, which leads to different mechanisms relative to a static matching model as in [Gayle and Shephard \(2019\)](#).

My research is further related to papers which have used relatively similar frameworks - dy-

namic bargaining models - to study different aspects of the tax and transfer system, such as for how long individuals are eligible to receive social assistance benefits (Low et al. (2018)), the EITC (Mazzocco, Ruiz, and Yamaguchi (2013)) or a comparison of joint and individual taxation (Bronson and Mazzocco (2019)). Relative to these papers, this paper adds the marriage market equilibrium, so that the distributions of potential partners are endogenously determined in equilibrium. Equilibrium models of marriage and intra-household allocations are still relatively rare (see e.g. Shephard (2019), Ciscato (2019), Goussé, Jacquemet, and Robin (2017), Chiappori, Dias, and Meghir (2018) or Reynoso (2019) for some recent work). Chade and Ventura (2002, 2005) study the impact of the differential tax treatment between singles and couples on the marriage market. In my analysis, the tax code is formally neutral with respect to marriage, due to individual taxation, but tax progressivity affects marriage rates by reducing wage inequality. The focus of this paper on intra-household bargaining is also related to Knowles (2012), who studies the role of bargaining for explaining the time trend in relative leisure between men and women. Finally, several papers have highlighted the role of the marriage market for inequality (e.g. Greenwood et al. (2016), Fernandez, Guner, and Knowles (2005) or Fernández and Rogerson (2001)). These papers abstract from bargaining and intra-household inequality, and thus focus on different aspects of the marriage market. In addition, they do not consider the impact of policies.

In the empirical literature, a growing number of studies has investigated the role of intra-household inequality. Lise and Yamada (2018) analyze Japanese panel data and relate allocations to differences in relative wages. Santaaulàlia-Llopis and Zheng (2017) study Chinese data and find that standard equivalence scales understate inequality in food consumption substantially and emphasize the role of luxury goods. Other papers have estimated collective household models to infer private consumption from micro-data on allocations (e.g. Lise and Seitz (2011), Cherchye, De Rock, and Vermeulen (2012); Cherchye et al. (2015, 2018)). Lise and Seitz (2011) conclude that intra-household inequality accounts for about 25% of consumption inequality in recent years and that it is crucial for the assessment of the time trend. In addition, empirical studies that test unitary and collective household models often reject unitary models (e.g. Attanasio and Lechene (2014), Lundberg, Pollak, and Wales (1997)) and find that intra-household allocations react to changes in outside options, which is consistent with models of bargaining.

The rest of this paper is organized as follows. Section 2 first shows stylized facts about intra-household allocations. Section 3 describes the model, which is calibrated in section 4. Section 5 shows the details of the policy simulations and discusses the results. Section 6 concludes.

2 Motivating facts

This section briefly describes stylized facts about intra-household allocations. The main data set being used is the "Time Use and Consumption" module of the Dutch LISS panel (see Cherchye, De Rock, and Vermeulen (2012) or Cherchye et al. (2017) for a detailed description). The module contains relatively detailed questions about household expenditures. Individuals are asked about

their private consumption as well as about overall spending for different categories and kids. The survey questions refer to personal expenditure for food, clothing, cigarettes, leisure time expenditure and a few other categories. At the same time, only a part of the overall consumption expenditure is directly assignable to a household member; the larger part is non-assignable (containing for example trips, housing expenditure or utility payments). I follow [Cherchye, De Rock, and Vermeulen \(2012\)](#) in their classification of non-assignable consumption. In practice, then, it is assumed that a fraction of non-assignable consumption is public.

In the following, I highlight several aspects which motivate some of the model assumptions. First, panel (a) of figure shows the ratio of the wife’s private (assignable) consumption relative to total private assignable consumption. This figure suggests that there are quite many couples in which spouses consume different amounts of private goods, which supports the notion of intra-household inequality in consumption.⁴ In a within-between decomposition of the variance of private consumption, the within-couple component has 51.1% the size of the between-couple component.⁵

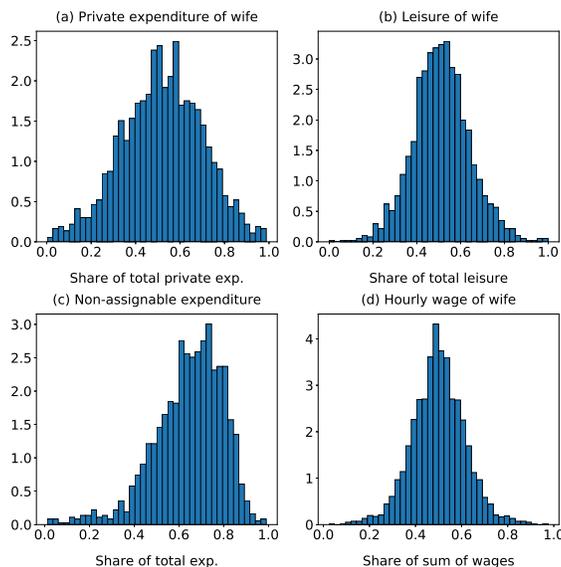


FIGURE 1

Notes: Panel (a) shows the share of private assignable consumption of the wife relative to total (assignable) private consumption in the household. Panel (b) and (d) do the same for hourly wage rate. Panel (c) depicts the share of expenditure used for non-assignable consumption. The data comes from the LISS panel.

Second, panel (b) shows the distribution of the leisure hours of the wife relative to total leisure of the household. The figure indicates that it is often the case that one spouse has more leisure

⁴Note that there are two potential sources of measurement error in consumption surveys. The first is an *infrequency bias*. The LISS asks individuals about consumption expenditure in a month. If, for example, one spouse makes a private purchase in one month and the other spouse purchases the same item next month, this would look like inequality in the survey. The second is a *recollection bias*: if spouses cannot perfectly estimate their personal consumption, this generates some noise.

⁵The within-between decompositions are described in more detail in section 5.2.

than their partner, which suggests that the allocation of leisure can be an important dimension of intra-household inequality.

Third, an important question regarding the potential for within-household inequality in consumption is the role of public goods. Panel (c) depicts the ratio of non-assignable expenditure, i.e. expenditure that cannot be attributed to the personal consumption of either spouse in the survey, relative to total expenditure and shows that it makes up for a relatively large fraction of consumption in many households (for example due to rent payments). This indicates that accounting for the publicness of some consumption is important.⁶

Finally, panel (d) shows the ratio of the hourly wage of the wife relative to the sum of wages. The figure is important since one would expect within-household inequality to be most relevant when spouses strongly differ in their earnings potential. Thus, the relative importance of within- and between household inequality should strongly depend on how assortative matching is in terms of wages. The figure provides an indication of how many couples there are in which either of the spouses has an economically stronger position. The focus is on *hourly wages*, rather than on income, since wages are a better indicator of the earnings *potential* of an individual, especially for married women. The data indicates that dispersion in wages is relatively large. In a within-between decomposition of hourly wages, the within-household component accounts for about 40% of the total variation. Interestingly, the distribution is relatively symmetric, in the sense that there are also many couples in which the wife has a higher wage than the husband.

Lise and Yamada (2018) report similar graphs for Japan. The main difference between the Netherlands and Japan is that the gender difference is much less pronounced in the Netherlands, whereas Japan has a large gender wage gap and the mean share of private consumption of women is considerably lower. In the Netherlands, by contrast, there is also a noticeable fraction of couples in which the woman has the higher wage rate. Earnings also exhibit a strong gender difference in the Netherlands, since part-time employment is very common for women.

3 Model

3.1 Overview

The model is an equilibrium search model of marriage, divorce, labor supply and savings. In each period, a new cohort of men and women is born and lives for a finite time horizon T . Each model period corresponds to two years and the terminal period is $T = 30$. Individuals participate in the labor market until retirement age $T_R = 21$. Individuals are heterogenous in their labor market ability and face shocks to their earnings potential over their lifetime. They decide on consumption and savings, as well as on splitting their time use between market work, home

⁶Note that classifying consumption as either private or public ultimately requires making assumptions about preferences. For example, one could regard a car as a public good, since it can to some extent be used by both spouses, or as a private good, if only one spouse cares about having a car. In addition, spouses could disagree over the attributes of the good, so that a given purchase benefits one spouse more than the other.

production and leisure. There are two consumption goods. One is private and the other one is public within the household and consumed equally by both spouses, in case an individual is married. The home good, which is produced using the time inputs of spouses, is also public within the household. The main features of the model are that it allows for within-household inequality in personal consumption and leisure, and that marriage and divorce decisions are endogenous.

Individuals start their lives single and can find a spouse on a search-based marriage market. In each period, singles may meet a potential partner from the distribution of available singles. Singles only meet potential partners who are from the same cohort and thus have the same age. At the time of a meeting, individuals take the economic characteristics of their potential partner and a non-economic 'love' component into account. Marriage takes place if both individuals prefer getting married over staying single for another period. Couples exogenously have children according to a stochastic process. The economic benefits from marriage are economies of scale in consumption and home production, since the public consumption good and home production time are public goods within the household, and the possibility to share labor market risk. However, not all goods are shared equally by both spouses. Personal consumption and leisure are consumed privately.

Decision-making in marriage is modeled as a limited commitment contract (Mazzocco, Ruiz, and Yamaguchi (2013)). This means that couples decide cooperatively on allocations. The Pareto weight of a spouse determines the weight in the objective function of the household and thus how favorable the allocation is towards this spouse. The Pareto weight is initially determined by Nash bargaining in the period when the couple gets married. The *outside option* in this bargaining situation is the value of staying single and meeting another potential partner in the next period. After marriage, the Pareto weight can be renegotiated when one partner would prefer to divorce given the current Pareto weight. The non-economic match quality of the couple can deteriorate over time, which, together with the wage shocks, may lead to divorce. Divorced individuals reenter the marriage market and can remarry in future periods.

The distributions of potential partners that singles can meet are endogenously determined in equilibrium. For each age group, these distributions depends on which individuals got married or divorced in previous periods and thus leave or reenter the pool of available singles. In a steady-state equilibrium, individuals have correct expectations over the future distributions of singles. Thus, finding an equilibrium corresponds to finding a fixed point between the expectations of agents and the implied distributions of singles.

In the following, the model will be described in more detail. Section 3.8 then shows illustrations about how intra-household allocations are determined and discusses how this is affected by the tax system.

3.2 Preferences and home production

Individuals differ in gender $g \in \{f, m\}$ and have preferences over consumption, leisure and a home good, which is produced with time as an input. There is a private and a public consump-

tion good and both goods are available to singles and married individuals. The public good C is consumed equally by both spouses when an individual is married, as opposed to the private consumption good c . The home good is denoted as D and is a public good within the household. Individuals have a total time budget of 1, which can be divided between market work (h_g), domestic work (d_g) and leisure (l_g):

$$h_g + d_g + l_g = 1$$

The per-period utility function from consumption, leisure and home production is given by:

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

Married individuals further get utility from the match quality of their marriage which will be discussed in more detail in section 3.5. The preference parameter for the home good D is allowed to depend on the presence of kids (b), which captures that couples with children typically spend more time on home production. The home good is produced according to the following home production technology, that takes the domestic work time of each spouse as the inputs.⁷

$$D = \begin{cases} \eta d_g & \text{when single} \\ (\eta_f(b)d_f^z + \eta_m(b)d_m^z)^{\frac{1}{z}} & \text{when married} \end{cases}$$

The parameter z controls the substitutability between the two time inputs. I assume that the productivity parameter of home time is the same for single men and single women, but allowed to differ between genders in marriage, depending on whether there are children ($\eta_g(b)$). When $\eta_f(b) > \eta_m(b)$, married women spend more time in home production than men do, even in couples where both partners have the same wage rate. This captures the empirical observation that women take a disproportionate share of total home production time. The specification could for example represent gender norms and does not necessarily reflect actual differences in 'productivity'.

The productivity parameter for married men is set to $\eta_m(b) = 1 - \eta_f(b)$. Individuals discount the future at rate β and have access to a risk-free asset that yields an exogenous interest rate R .

3.3 The labor market and the government

Labor supply choices belong to a discrete set of hours $\mathbb{H} = \{h_1, \dots, h_k\}$. The (log) wage of an individual is given by:

$$\log w_{i,t} = \kappa_i + \epsilon_{i,t} + \kappa_g^G + u_{it}$$

⁷The advantage of this specification of preferences and home production is that it allows for closed form solutions for consumption (c_f , c_m and C) conditional on total household expenditure and the Pareto weight, which makes the model more tractable. It would be an interesting extension to consider a home production function where the public good enters as a substitute to time (as in Knowles (2012) or Lise and Yamada (2018)). For example, the interpretation could be that households can partially outsource their domestic work. However, it would increase the computational burden. Similarly, allowing for non-separability between consumption and leisure would be an interesting extension.

κ_i represents the ability of the individual. κ_g^G refers to the gender wage gap.⁸ $\epsilon_{i,t}$ is a Markov chain that represents *persistent* changes in earnings potential over the life-cycle. u_{it} is a *temporary* (i.i.d.) shock and follows a normal distribution with mean 0 and variance σ_T .⁹ In retirement, individuals get a fixed replacement rate of their full-time earnings.

The government collects revenue through a tax on labor income. The tax system is represented by a parametric tax function, using the functional form from [Heathcote, Storesletten, and Violante \(2017\)](#):

$$T(y) = \max \left((1 - \psi_1 y^{-\psi_2}) \cdot y, 0 \right)$$

This tax function implies that the after-tax income (conditional on being positive) is given by:

$$y - T(y) = \psi_1 y^{1-\psi_2}$$

The parameter ψ_1 determines the level of the tax, since it proportionally shifts the after-tax income. ψ_2 determines the progressivity of the tax system. Progressivity can be defined as the average tax rate ($\frac{T(y)}{y}$) being increasing in income and is commonly measured by the progressivity tax wedge (see e.g. [Holter, Krueger, and Stepanchuk \(2019\)](#) for details):

$$PW(y_1, y_2) = 1 - \frac{1 - T'(y_2)}{1 - T'(y_1)} = 1 - \left(\frac{y_1}{y_2} \right)^{\psi_2}$$

Thus, given the functional form for the tax function, the progressivity wedge only depends on ψ_2 . Married couples are taxed individually in the Netherlands, so that the tax liability is calculated for each spouse separately.¹⁰ For simplicity, this specification abstracts from aspects like transferable deductions between spouses, which would introduce a small degree of dependence between partners.¹¹ Note that the tax function $T(y)$ is restricted to be non-negative. This has the interpretation that $T(y)$ models only the tax and not the transfer system.¹²

Transfers are modeled as a simple income floor depending on *family income*, which prevents households from experiencing very low consumption levels. Households with an income below B_{min} receive a transfer to reach this level, that is financed out of tax revenue. The budget constraint of the government equates the expected life-time revenue from a newly born cohort to an

⁸Note that the model abstracts from human capital accumulation/learning by doing. This was included in an earlier version of the paper, but replaced by a richer distribution of permanent types to model the income distribution and assortative mating in a more realistic way, which is important for the study of inequality. The current specification can be thought of as frontloading the earnings potential of an individual into the ability type. Adding learning-by-doing on top of the heterogeneity and persistent earnings risk would lead to a very large state space.

⁹The distinction between temporary and transitory shocks is often made in the literature (see e.g. [Krueger and Perri \(2006\)](#)). In addition, the temporary shock is helpful for computational reasons, since it makes the problem more smooth.

¹⁰The actual Dutch tax schedule consists of several steps, within which the marginal tax rate is constant. Approximating the schedule with a tax function with few parameters makes policy experiments more straightforward. In the calibration section, I will discuss the choice of parameters and how it compares to the actual schedule in more detail.

¹¹See [de Boer et al. \(2018\)](#) for a detailed overview of the institutional setting in the Netherlands.

¹²Excluding transfers from $T(y)$ is required since $T(y)$ is applied to the *individual* labor income of spouses, whereas transfer are typically computed on the basis of *family* income.

exogenous spending requirement \bar{G} :

$$\text{Rev}(\psi_1, \psi_2) = \mathbb{E} \left[\sum_t \beta^t \left((T(y_{i,t}) - B(y_{i,t})) \right) \right] = \bar{G}$$

Here, B indicates whether an individual gets the social assistance benefit. The revenue \bar{G} is used to finance non-valued government expenditure.¹³

Note that the model features the standard policy trade-off of progressive taxation (as in e.g. [Conesa and Krueger \(2006\)](#) or [Heathcote, Storesletten, and Violante \(2017\)](#)). An increase in progressivity (ψ_2) is potentially welfare-improving because it allows lowering tax rates for lower incomes (through ψ_1) while keeping the government budget constant. On the one hand, this can be desirable since it redistributes to low-income individuals who have a higher marginal utility from income. In addition, a budget-balanced increase in progressivity provides more insurance against wage fluctuations. On the other hand, the increase in progressivity can also reduce welfare because of the labor supply disincentives through high marginal tax rates.

3.4 Singles

Singles decide on consumption, savings and the time spent for home production, leisure and labor supply. They are characterized by their gender, ability, the productivity shock, the presence of children, if there are any from a previous marriage, and assets. Thus, the state vector of a single is:

$$\omega_{i,t}^S = (g, \alpha_{i,t}, \epsilon_{i,t}, b_{i,t}, A_{i,t})$$

In the beginning of each period, individuals observe the realization of the persistent productivity shock and whether the children grow up. Their decision problem is described by the following value function and budget constraint:

$$\begin{aligned} V_{i,t}^S &= \max_{c, C, h, l, d, A'} \{u(c, C, l, D) + \beta \text{EV}_{i,t}^S(A')\} \\ c + C + A' &= w_{i,t}h - T(w_{i,t}h) + RA_{i,t} \\ d + h + l &= 1 \end{aligned}$$

The current asset level is denoted as $A_{i,t}$, while A' is the amount of savings for the next period. The continuation value $\text{EV}_{i,t}^S(A')$ includes the expected utility from being single in the future, as well as from future marriage, divorce and remarriage. State variables, or objects that depend on state variables, are indexed by the individual and period.

In addition to their ability, the initial condition of singles also contains a realization of the productivity shock. To ensure that the initial draw is consistent with the earnings process, it is obtained by drawing from the distribution of the process after two periods, starting with the

¹³This can also be interpreted as the government producing an economy-wide public good P , which provides each individual with an additive utility $u^P(P)$ and does not affect choices. An alternative assumption would be to rebate the tax revenue to workers, which would provide an additional minimum income.

median wage shock. In addition, individuals draw a small initial amount of assets, that could e.g. be interpreted as a parental transfer. The assets are drawn from a uniform distribution between 0 and 10% of the full-time earnings of the individuals.¹⁴

3.5 Couples

Before turning to matching and marriage decisions, it is useful to discuss the decision problem of married couples and, in particular, the role of the Pareto weight in more detail. Married couples decide on consumption, labor supply and home production time. The utility of a married individual ($g \in \{f, m\}$), conditional on these choices is:

$$U_{i,t}^g(c, C, l, D, A') = u(c, C, l, D) + \theta_{i,t} + \beta \text{EU}_{i,t}^g(\lambda_{i,t}, A')$$

On top of the economic utility $u(c, C, l, D)$, married individuals also experience the additive 'love shock' $\theta_{i,t}$, which represents the non-economic quality of the current relationship. Match quality evolves stochastically according to an $AR(1)$ process with persistence ρ_θ and variance σ_θ :

$$\theta_{i,t} = \rho_\theta \theta_{i,t-1} + \epsilon_{i,t}^\theta + \bar{\theta}_{\alpha_f, \alpha_m}$$

$\epsilon_{i,t}^\theta$ represents the innovation of the match quality process.¹⁵ $\bar{\theta}_{\alpha_i, \alpha_{j(i)}}$ is a preference term that generates additional utility depending on the combination of ability types (α_f and α_m) of the couple. In practice, I assume that

$$\bar{\theta}_{\alpha_f, \alpha_m} = \begin{cases} \bar{\theta} & \alpha_f \text{ and } \alpha_m \text{ in same ability quartile} \\ 0 & \text{else} \end{cases}$$

The interpretation of this specification is that individuals enjoy additional utility when they are married to someone with a broadly similar ability type, for example because of shared interests or social conventions. Thus, $\bar{\theta}$ will be referred to as the homophily parameter. This helps the model to generate a realistic degree of assortative mating.¹⁶

$\text{EU}_{i,t}(\lambda_{i,t}, A')$ is the continuation value of marriage that includes the possibility of divorce in the next period. Married couples exogenously have children with a probability $p^b(t)$, that declines with age and is zero once the couple reaches the age of 40. Having children is a binary state ($b \in \{0, 1\}$), which can be thought of as all couples having two kids. Children leave the household with probability $p^{b,g}$.¹⁷ The state space of married couples includes the current Pareto weight λ , the ability level (α), and productivity shock (ϵ) of each spouse as well as the assets, the current

¹⁴The initial asset endowment is useful for computational reasons, since it makes the (equilibrium) asset distributions in the first periods more smooth.

¹⁵In practice, the match quality process is approximated with a discrete process using the Rouwenhorst method.

¹⁶This is comparable to Greenwood et al. (2016), who have a similar preference term for college and non-college individuals. In my context, some alternative specifications would be possible, for example to let the preference term depend more continuously on the type difference, for example as some function of $|\alpha_f - \alpha_m|$.

¹⁷This assumption avoids keeping track of the age of kids, which would greatly enlarge the state space.

level of match quality (θ) and the presence of children (b):

$$\Omega^C = \{(\lambda, \alpha_f, \alpha_m, \epsilon_f, \epsilon_m, A, \theta, b)\}$$

The most important variable that determines the allocation within the couple is the Pareto weight $\lambda_{i,t}$. The weights are normalized such that $\lambda_{i,t}$ is the weight of the wife and $1 - \lambda_{i,t}$ is the weight of the husband. The Pareto weight is endogenously determined. For the moment, focus on a couple with a given weight Pareto weight $\lambda_{i,t}$, which was obtained in a previous period. The household problem is to maximize the weighted sum of utilities of wife and husband, subject to the budget and time constraints:

$$\begin{aligned} \max_{\mathcal{C}=(c_f, c_m, C, A', h_f, h_m, d_f, d_m, l_f, l_m)} \quad & \lambda_{i,t} U_{i,t}^f(\mathcal{C}) + (1 - \lambda_{i,t}) U_{i,t}^m(\mathcal{C}) \\ c_f + c_m + C + A' = & w_{f,t} h_f + w_{m,t} h_m - T(w_{f,t} h_f) - T(w_{m,t} h_m) + R A_{i,t} \\ h_f + d_f + l_f = & 1 \\ h_m + d_m + l_m = & 1 \end{aligned}$$

Thus, a high value of $\lambda_{i,t}$ corresponds to a high weight placed on the utility of the wife and a large amount of private consumption and leisure for her, whereas a low value of $\lambda_{i,t}$ leads to high utility for the husband. The utility levels evaluated for the optimal choices are denoted as $U_{f,t}^*(\lambda_{i,t}, \omega_{i,t}^M)$ and $U_{m,t}^*(\lambda_{i,t}, \omega_{i,t}^M)$, given λ and the state variables of the couple ω_c .

Since spouses can unilaterally file for divorce in the beginning of a period, allocations have to make sure each spouse prefers staying married over divorcing. If this is not the case, the Pareto weight can be adjusted to ensure that the spouse who would want to divorce stays in marriage (limited commitment).¹⁸ The participation constraint requires that each spouse is better off in marriage than as a single:

$$U_{g,t}^*(\lambda_{i,t}, \omega_c) \geq V_{g,t}^S(A_{i,t}/2, \alpha_{g,t}, \epsilon_{g,t}, b_{i,t})$$

Assets are split equally in divorce and children stay with the mother, who also takes them into the next marriage.¹⁹ The Pareto weight of the couple stays constant until a participation constraint is violated. In case one spouse would want to divorce, it may be possible to increase their Pareto weight to make them stay in marriage. If this is feasible and the other spouse also still prefers staying in marriage, the Pareto weight is adjusted in favor of the spouse who wants to leave otherwise. Efficiency requires making the *minimal* adjustment relative to the previous Pareto weight,

¹⁸See Chiappori and Mazzocco (2017) for a theoretical overview of limited commitment models.

¹⁹The equal division of assets is a simplification that is often made in the literature (see e.g. Fernandez and Wong (2017); Mazzocco, Ruiz, and Yamaguchi (2013)) and close to the legal regime in the Netherlands. Importantly, this assumption rules out that couples can write prenuptial contracts to e.g. keep separate property, which is rare in practice. Note that allowing couples to keep separate property would *increase* the importance of bargaining, since then the spouse with the higher bargaining power would be able to keep more property in their name, increasing the share of assets in divorce.

so that the spouse is just indifferent between divorcing and staying in the marriage. If such an adjustment is not possible, divorce occurs.

The interpretation of limited commitment is that spouses behave as under a constrained efficient contract, given the constraint that the contract can be renegotiated if the participation constraint of a spouse is violated in a future period. As a result, risk-sharing in marriage is imperfect. Shocks to wages and match quality can trigger changes in the Pareto weight of the couple. This is the more likely the lower the match quality of the couple, since rebargaining requires that one individual would prefer to get divorced given the current bargaining weight, which is less likely when match quality is high.²⁰

3.6 Meetings and marriage

Singles participate in the marriage market with probability $m(t)$. This probability declines with age to capture that it is harder to meet potential partners later in life, when fewer people are single.²¹ Singles who enter the marriage market randomly meet a potential partner from the (endogenous) pool of available singles. For tractability, I assume that singles only meet others from the same cohort, so that spouses have the same age.²² The characteristics of the potential spouse are their ability, productivity, assets and children. The probability of meeting a spouse with a state vector (ϵ, α, A, b) is given by:

$$M_{t,g}(\epsilon, \alpha, A, b) = m(t) \cdot \Lambda_{t,g}(\epsilon, \alpha, A, b)$$

$\Lambda_{t,g}$ reflects the likelihood of meeting a potential spouse with certain characteristics, given the distributions of singles at each age. $\Lambda_{t,g}$ is endogenously determined in equilibrium, since it depends on which individuals got married or divorced at previous younger ages and on how much they saved and worked. The characteristics of each potential spouse are denoted as $\omega_{i,t}$. In addition, an initial realization of match quality $\theta_{i,t}$ is drawn from a $N(\mu_{I,\sigma}, \sigma_{I,\theta})$ distribution for each match, which captures the non-economic quality of the potential marriage.

At the time of a meeting, have to decide whether to get married and, if so, set an initial value for the Pareto weight $\lambda_{i,t}$. They observe the match quality draw and the persistent characteristics of the potential partner. For technical reasons, it is helpful to assume that the temporary wage shocks are realized *after* individuals have decided on whether to get married. This ensures that the *expected* utility from marriage $\mathbb{E}_w V_{g,t}^M(\lambda, \theta_{i,t}, \omega_{f,t}, \omega_{m,t})$ is a smooth function of the Pareto weight, even though labor supply choices are discrete.²³ The expectation \mathbb{E}_w refers to the expectation over the temporary wage shocks. Note that the temporary shocks are i.i.d., so that the timing assumption should have a small impact on marriage decisions or bargaining. Importantly, individuals

²⁰Evidence from [Lise and Yamada \(2018\)](#) supports these models of decision-making, as they find that decision power changes infrequently and more often before divorce and for big shocks.

²¹A richer specification would be to model $M(t)$ endogenously as a function of the fraction of singles at each age group, as in [Guvenen and Rendall \(2015\)](#).

²²Otherwise, one would need to keep track of the age of both spouses, which would greatly enlarge the state space.

²³The smoothness of the Pareto frontier is helpful when computing the bargaining solutions.

observe the *persistent* wage shock of the potential partner before making the marriage decision.

When deciding to get married, individuals compare the value from staying single for one more period and drawing another potential partner in the next period ($\mathbb{E}_w V_{g,t}^S$) to the value of getting married to the current potential partner ($\mathbb{E}_w V_{g,t}^M$). Given a value for the Pareto weight λ , the *gain from marriage* of individual g is:

$$G_g(\lambda, \omega_{f,t}, \omega_{m,t}, \theta_{i,t}) = \mathbb{E}_w V_{g,t}^M(\lambda, \omega_{f,t}, \omega_{m,t}, \theta_{i,t}) - \mathbb{E}_w V_{g,t}^S(\omega_{g,t})$$

For marriage to take place, there must be a Pareto weight $\lambda \in (0, 1)$ such that both individuals prefer marriage over staying single:

$$G_g(\lambda, \omega_{f,t}, \omega_{m,t}) \geq 0 \quad \forall g \in \{f, m\}$$

If there is no Pareto weight such that both individuals prefer marriage, the match is rejected and the individuals stay single for another period. In other cases, there is an interval of weights $[\lambda, \bar{\lambda}] \subset (0, 1)$ for which both individuals prefer marriage over staying single. The Pareto weight also has the interpretation of a transfer, since spouses can reduce their own utility in order to increase the utility of their partner.²⁴ The Pareto weight is determined by Nash bargaining, which chooses the weight λ which maximizes the Nash product:²⁵

$$\tilde{\lambda} = \operatorname{argmax}_{\lambda} G_f(\lambda, \omega_{f,t}, \omega_{m,t}, \theta_{i,t}) \cdot G_m(\lambda, \omega_{f,t}, \omega_{m,t}, \theta_{i,t})$$

Note that this is the most essential model ingredient that determines intra-household allocations. Section 3.8 provides a more intuitive discussion of how Pareto weights are determined and how this can be influenced by the tax system. Nash bargaining relates the allocation in marriage to the relative 'outside options' of each individual. The outside option is the utility of staying single for another period, as previously defined:

$$V_{i,t}^S = \max_{c, C, h, l, d, A'} \{u(c, C, l, D) + \beta \mathbb{E} V_{i,t}^S(A')\}$$

With the notation from this section, we can state the continuation value of singlehood as the expected utility from future periods of singlehood and the expectation over all potential marriages:

$$\mathbb{E} V_{g,t}^S(A') = \int_{\omega^M = (\omega_{f,t+1}, \omega_{m,t+1}, \theta)} M(\omega_{t+1}^M) \cdot \mathbb{E}_w V_{g,t+1}^m(\lambda(\omega_{t+1}^M)) + (1 - M(\omega_{t+1}^M)) \cdot \mathbb{E}_w V_{g,t+1}^S \, dF(\omega_{t+1}^M)$$

ω_t^M is a vector containing the characteristics of both individuals and the match quality realization. $M(\omega_t^M)$ is an indicator variable that takes the value of 1 if marriage occurs given these match char-

²⁴In technical terms, utility is *imperfectly transferable*, since giving up one util does not increase the utility of the partner by an equal amount. See Galichon, Kominers, and Weber (2018) for a theoretical discussion.

²⁵Note that the general formulation of Nash bargaining contains a parameter for 'bargaining strength', which corresponds to the exponents of the Nash product. Here, it is assumed that this parameter is 0.5.

acteristics and $\lambda(\omega_t^M)$ is the Pareto weight that is chosen in this case. In particular, the expression for $EV_{g,t}^S(A')$ highlights that outside options are determined both by the living standard while single and by the full distribution of potential marriages in the future.

3.7 Marriage market equilibrium

3.7.1 Definition of equilibrium

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

1. the policy functions $(A, c, C, h, l, d) = P_{g,t}^S(\omega^S)$ solve the problem of singles
2. the policy functions $(A, c_f, c_m, C, h_f, h_m, l_f, l_m, d_f, d_m) = P_t^M(\omega^M)$ 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 getting married and λ the initial Pareto weight
5. the implied distributions of singles, $\Lambda_{t,g}(\epsilon, H, \alpha, A)$, are consistent with the distributions that are used to determine the optimal choices and value functions from (1) - (4)

3.7.2 Discussion and computation

The main concern of the marriage market equilibrium is to find distributions of singles of each gender and age group $(\Lambda_{t,g}(\epsilon, \alpha, A, b))$. Since these distributions determine which potential partners an individual can meet at each age, they partly determine the forward-looking decisions about marriage, divorce and savings. At the same time, these decisions feed back into the distributions of available singles at each age. Thus, equilibria need to be computed via fixed-point iteration. Starting with a guess for the distributions of singles at each age, one can solve the life-cycle problem recursively to determine the value functions. Then, the actual distributions of singles, given the guess of the distribution, are computed from a simulated panel of 1 million individuals and the guess is updated. This is repeated until convergence. Since there are relatively many state and choice variables, computing the equilibrium is fairly time-consuming - the parametrization used for the policy experiments runs on a cluster using 448 cores. The computational details are described in appendix A.

3.8 How does tax progressivity affect intra-household inequality?

To illustrate the mechanisms of the model, this section discusses an example to show how Pareto weights are determined and highlights the aspects that are relevant for the effects of progressivity. First, consider a simple static example, in which the economy consists of two individ-

uals. Individuals consume the public and the private good and labor supply is fixed. The utility function over private and public consumption is given by:

$$u(c, C) = \alpha_c \frac{c^{1-\gamma}}{1-\gamma} + (1 - \alpha_c) \frac{C^{1-\gamma}}{1-\gamma}$$

Figure 2 illustrates the bargaining situation. The individuals need to determine the weight of the woman λ , the weight of the man being $1 - \lambda$. The decision weight determines consumption and leisure in marriage. Each value of λ corresponds to a point on the Pareto frontier, which is a combination of utility levels for the individuals. The higher λ , the higher the utility level of the woman. The black (vertical) line is the value of singlehood for the woman and the green (horizontal) line is the one for the man. In the graph on the left side, wages are unequal and the value of being single is lower for the woman than for the man. The final allocation must lead to each utility level being higher than the value of singlehood. This is the case for a range of values of the Pareto weight, which corresponds to the line segment between the points where the Pareto frontier intersects the values of singlehood. Whenever this is the case, marriage takes place. The bargaining solution picks one particular point on the Pareto frontier, which is marked by the dot. In the case of unequal wages, the Pareto weight is more favorable for the man, who receives more consumption, reflecting the difference in outside options.

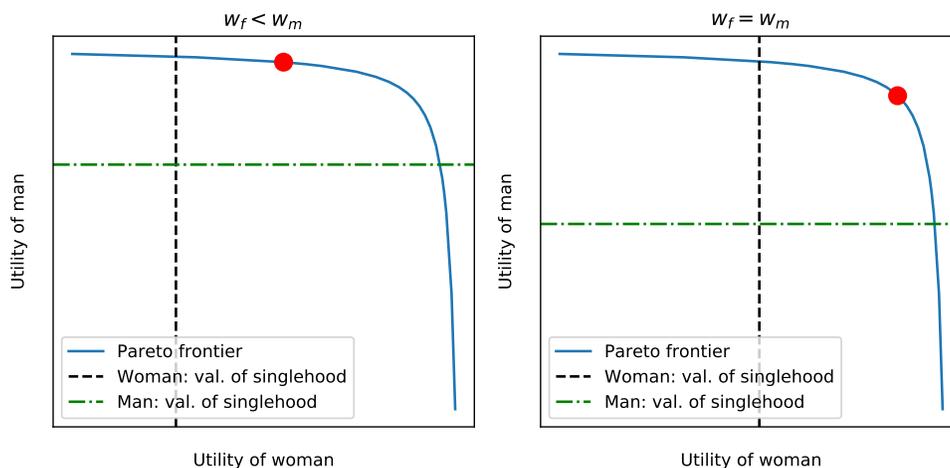


FIGURE 2: Illustration - Determination of Pareto weight

Notes: The graphs illustrate the impact of changing relative wages on the bargaining outcome in a simple example. The blue line is the Pareto frontier. Each point on the frontier corresponds to a Pareto weight. The dot shows the bargaining outcome given the outside options.

The graph on the right side shows the case where wages are equal. As a result, the value of being single is now equal for the two individuals. The new allocation is marked by the dot and assigns equal utility in marriage to each partner. In this example, total household income and thus the Pareto frontier were unchanged. Such bargaining effects are in line with a range

of empirical studies that tests unitary and collective household models and finds that allocations react to changes in outside options (e.g. [Attanasio and Lechene \(2014\)](#), [Lundberg, Pollak, and Wales \(1997\)](#)).

In addition to the outside options, intra-household allocations are also determined by the shape and location of the Pareto frontier (which is unchanged in the example in figure 2). With Nash bargaining, the initial Pareto weight is chosen such that the product of the gains from marriage is maximized.²⁶ The main determinants of the Pareto weight are best illustrated by the first-order condition of the bargaining solution:

$$\frac{V_{f,t}^M(\lambda) - V_{f,t}^S}{V_{m,t}^M(\lambda) - V_{m,t}^S} = \frac{\partial V_{m,t}^M}{\partial V_{f,t}^M}(\lambda)$$

The terms $V_{g,t}^M(\lambda) - V_{g,t}^S$ represent the gains from marriage for each individual $g \in \{f, m\}$, given a Pareto weight λ . Allocations are determined by the relative gains from marriage of each individual and the slope of the Pareto frontier, which indicates how transferable utility is between spouses. When both individuals have a relatively high gain from marriage irrespective of the Pareto weight, which is for example the case when the match quality of the marriage ($\theta_{i,t}$) is high, the left-hand side of the equation is close to unity. This shifts the Pareto weight towards 0.5, where the slope of the frontier is also close to unity. Besides public consumption, home production and match quality, the gains from marriage also include risk-sharing in this model, which affects the Pareto weight.

Influencing the intra-household decision weight can improve utilitarian welfare, since the gain of one person usually offsets the loss of the other (utility is imperfectly transferable). In the example discussed above, equalizing wages, such that the decision weight is set to 0.5 is the optimal policy for a utilitarian government that places equal weights on both individuals.

The discussion so far focused on static situations, in which the outside option is to stay single for the period. In the full model, the outside option for each individual is not to stay single for the rest of the lifetime, but includes the possibility of future (re-)marriage. As a result, outside options depend on the distributions of potential partners that will be available in the future, the probability of getting married in the future, and future Pareto weights. Progressivity affects outside options by equalizing living standards while single and also by changing future marriage market outcomes. Note that the marriage market equilibrium implies that the distributions of potential partners that individuals meet are endogenous and changes due to policy reforms. Changes in progressivity can

²⁶Note that there are other bargaining solutions. A common alternative to Nash bargaining is the ‘Egalitarian’ bargaining solution (see [Kalai \(1977\)](#) or [Knowles \(2012\)](#)), which chooses the Pareto weight to equalize the gains from marriage (the surplus) of each spouse. With linear utility, Nash and Kalai bargaining are equivalent. With concave utilities, Nash bargaining also takes the slope of the Pareto frontier into account. The main difference is that Kalai bargaining leads to a higher elasticity of the Pareto weight with respect to the utilities. Thus, with Kalai bargaining, there would be more intra-household inequality and a stronger reaction of Pareto weights to the tax system. As discussed in section 4.2, the implied elasticities of Nash bargaining seem quite in line with empirical evidence. A second difference is that utility terms which are shared equally by both spouses (such as the non-economic ‘love’ component $\theta_{i,t}$) do not influence the Kalai bargaining outcome, while the Nash outcome does depend on $\theta_{i,t}$.

lead to changes in selection into marriage, which is reflected by the equilibrium distributions of types. In addition, progressivity changes how much different individuals can save and therefore the asset distributions of singles.

4 Calibration

4.1 Calibrated parameters

A number of parameters are set before the calibration routine. The coefficient of risk aversion is set to 1.5, which is a standard value in the literature. The discount factor is 0.9409, which implies a yearly discount factor of 0.97 (one period are two years). The interest rate is set to $R = 1/\beta = 1.063$. To set the preference parameter for the private consumption good, α_c , I use the fact that the preference specification implies a closed form solution for the share of public consumption of total expenditure, which only depends on the Pareto weight and the coefficient of risk aversion. Setting the preference parameter α_c to 0.277 implies that households spend roughly 61% of their expenditure on public goods.²⁷ Since wages are exogenous in the model, the wage process can also be set externally. The parameters that need to be set are the variance and persistence of the wage shock. This is done using data from the Dutch Socioeconomic Panel, due to its larger sample size. The wage process and the ability distribution are calibrated to match the cross-sectional dispersion of wages, as well as within-individual wage changes. The fit of the wage moments and the corresponding parameters are shown in the appendix. The remaining parameters are calibrated to match a set of data moments. The calibrated parameters are the preference parameters for leisure and the home good, the variance and persistence of the match quality shock, the meeting rate in the last period and the homophily parameter for assortative matching.²⁸ Regarding the moments, I match the share of individuals that is currently married or cohabiting at age 20 and 36, the share that ever married or cohabiting by age 36, the share that ever experienced a divorce or separation by ages 36 and 42.²⁹ These moments are constructed based on the Dutch Kinship Survey, which provides retrospective data on life histories. In addition, I match women's average hours spent on domestic and market work. These moments are based on the LISS panel and included separately by the number of kids, since the preference for home production is allowed to differ for couples with kids. Finally, to get a realistic degree of assortative matching, I also target the share of the total variance of hourly wages that is due to the variation within households. This moment

²⁷The exact share varies a little depending on the Pareto weight. The motivation for the data target is the procedure described in [Cherchye et al. \(2017\)](#): since only a part of the consumption expenditure is assignable in the LISS data, the authors assume that 50% of the remaining expenditure is public and the rest is private. With this assumption, I obtain 61% as the mean share of expenditure on public goods.

²⁸Recall that the meeting rate is assumed to be linear, so that only the rate in the last period needs to be determined. The meeting rate in the first period is set to 1. The decline in the meeting rate helps the model to rationalize that the fraction of married individuals is lower than one at older ages.

²⁹In the following, the term 'married' will refer to both married and cohabiting couples in the context of the data. Cohabitation is included in the data moments to accurately target the share of individuals in long-term partnerships, which would otherwise be understated. Note that there are some legal differences between married and cohabiting couples that are not captured by the model (relating e.g. to the division of assets upon divorce).

is primarily matched by the preference term for homophily. Intuitively, if couples were randomly matched, the share of the within-variation would be high since there would be many couples with unequal wages. Interestingly, this measure is relatively high in the data (0.4), suggesting that wage variation within couples is quite substantial.

The model is calibrated by minimizing the distance between model and data moments. Solving for the equilibrium distributions for each trial set of parameters would be computationally very expensive. Thus, I experimented with different approaches to reduce this computational burden. In practice, I start with a reasonable initial guess for the distributions, solve and simulate the model only once for each trial parameter vector, then update the distributions with a very small weight on the new distributions and proceed with the next trial parameter vector. In particular, this avoids solving for the full distributions for parameter vectors which would deliver a bad fit to the data anyway. After many iterations of the optimization algorithm, the distributions are close to the equilibrium distributions. The final parameters and distributions from this procedure are then used as starting points for the usual fixed point iteration, which converges after a few steps and delivers a good fit to the data. The calibrated parameter values are shown in table 1. Table 2 compares the model moments to the data moments that are targeted in the calibration.

TABLE 1: Calibrated parameters

Parameter	Value
Scale of leisure, women (α_l)	1.20
Scale of home production, no kids ($\alpha_D(0)$)	0.21
Scale of home production, kids ($\alpha_D(1)$)	0.68
Home productivity, women (η_f)	0.62
Homophily term ($\bar{\theta}$)	0.73
Variance of match quality, married (σ_l)	1.29
Autocorrelation of match quality (ρ_l)	0.89
Mean of initial match quality ($\mu_{I,\theta}$)	-0.21
Variance of initial match quality ($\sigma_{I,\theta}$)	0.74
Meeting rate in last period	0.01

TABLE 2: Model fit

	Model	Data
Work hours, women, couples without kids	0.37	0.35
Work hours, women, couples with kids	0.20	0.20
Home hours, women, couples without kids	0.17	0.21
Home hours, women, couples with kids	0.42	0.48
Leisure, women, couples without kids	0.46	0.44
Leisure, women, couples with kids	0.39	0.32
Mean share of housework (women)	0.61	0.60
Currently married, age 20	0.25	0.24
Currently married, age 36	0.82	0.79
Ever married, age 36	0.93	0.92
Ever divorced, age 36	0.25	0.31
Ever divorced, age 42	0.32	0.31
Share of within-couple wage variance	0.37	0.40

Notes: This table summarizes the fit of the model. The time-use moments are expressed as the fraction of total non-sleeping time which are used for the corresponding activity.

4.2 Model implications and comparison to data

Before turning to the results, this section first shows how the model compares to the data, in terms of inequality within and between households. Figure 3 shows the relative private consumption and leisure within couples, expressed as the ratio of female consumption and leisure relative to the sum of both household members. Thus, a value of 0.5 corresponds to both spouses having equal amounts. Relative consumption and leisure are untargeted in the calibration and are determined by the bargaining solution and the marriage market equilibrium. Both in the model and in the data, there is more dispersion in consumption than in leisure. In the model, this is driven by the fact that high-wage individuals have high bargaining power vis-a-vis low-wage individuals, but at the same time, the ‘price’ of their leisure is higher due to the higher opportunity cost. This effect dampens the dispersion in relative leisure.³⁰ In addition, the figure shows that the model generates less dispersion in consumption and leisure than present in the data, leaving room for alternative explanations of intra-household inequality (other than bargaining) that are not captured by the model, such as preference heterogeneity, which could be introduced as a residual.³¹

³⁰In the data, an alternative explanation for the difference in dispersion could also be that leisure can likely be measured more accurately than consumption.

³¹Preference heterogeneity would be an interesting extension, partly because it could *increase* the potential for intra-household inequality: when the preferences of spouses differ, bargaining power becomes more important.

To assess whether the bargaining solution gives empirically plausible outcomes, one can compare the model-implied elasticity of the Pareto weight with respect to relative wages to empirical estimates. Lise and Yamada (2018) estimate that on average increasing the difference in wages at the time of marriage by 10% translates into a 2.3% difference in the Pareto weight. In the model, the corresponding change in the Pareto weight is 1.9%, indicating that the implied elasticity of the Pareto weight seems reasonable and compares well to empirical studies.³² In particular, this elasticity of the Pareto weight is an important statistic for policy analysis. Finally, the figure also contains the distribution of household income of couples in the model and in the data, indicating that the model does a reasonable job at capturing inequality across households.

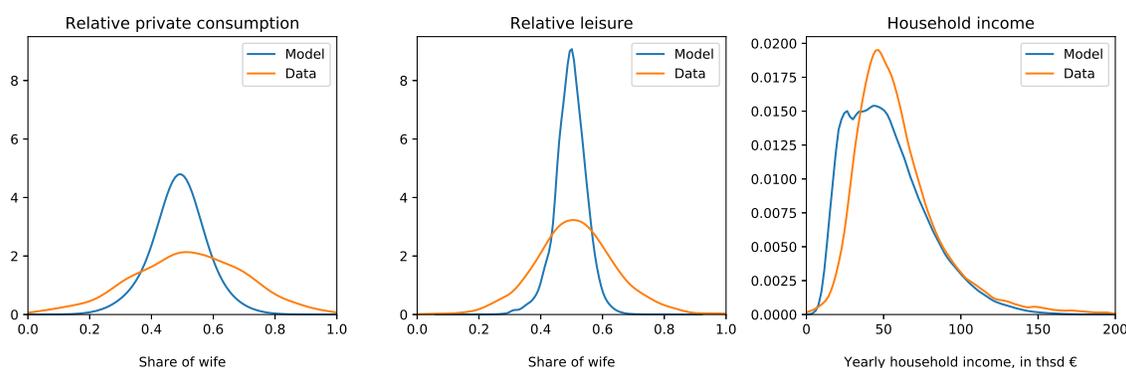


FIGURE 3: Inequality within and across couples

Notes: The figure shows the distribution of relative (private) consumption and leisure within households and the corresponding distributions in the data. Relative allocations are computed as the amount of the wife relative to the sum within the household (for example, $\frac{c_f}{c_f+c_m}$). In addition, the figure shows the distribution of (gross) household income.

To illustrate the Pareto weights underlying the intra-household allocations, table 3 shows the mean Pareto weight for each combination of ability group. Each group contains three ability types. The Pareto weights vary substantially across couples. In cases in which husband and wife are in the same ability group, the weights are relatively even on average. In unequal marriages, the individual with the better outside option gets a higher share of the surplus. For example, in couples which the wife is in the lowest ability group (L) and the husband in the highest (H), the mean Pareto weight is 0.26.

³²Note that the model-implied elasticity of the Pareto weight could be increased by using Kalai bargaining instead of Nash bargaining. However, in that case, the elasticity would be higher than the value found by Lise and Yamada (2018).

TABLE 3: Pareto weights by ability group of wife and husband

	Men - L	Men - LM	Men - UM	Men - H
Women - L	0.51	0.44	0.36	0.26
Women - LM	0.58	0.50	0.43	0.32
Women - UM	0.65	0.59	0.49	0.39
Women - H	0.74	0.68	0.60	0.48

Notes: The table shows the mean Pareto weights of couples across ability groups. "L" is the lowest group, "H" the highest. For example, group "L" contains types 1, 2 and 3.

5 Results

In this section, I first use the calibrated model to analyze inequality given the current tax schedule in more detail. How much of inequality is due to inequality *within* couples and *across* couples? How much is due to singles and due to the difference in means between singles and married individuals? I then describe the policy experiment - an increase in tax progressivity - and show how the reform affects inequality through its impact on within-family allocations, marriage and divorce. To complement the analysis of (cross-sectional) *inequality*, I finally also turn to the effect of the reform on (ex-ante) *welfare*.

5.1 Inequality Of What?

The model allows to go a step further than studying income inequality and consider inequality in consumption and utility more generally. Consumption inequality has been widely studied (see for instance [Krueger and Perri \(2006\)](#), [Blundell and Preston \(1998\)](#), [Blundell, Pistaferri, and Preston \(2008\)](#)). As demonstrated by [Lise and Seitz \(2011\)](#), consumption inequality is potentially understated by ignoring within-household inequality and assuming that consumption is shared equally within the household instead.

To get a complete picture of inequality, one would ultimately be interested in comparing *utility* in the population, since this captures that individuals differ not only in consumption, but also in their leisure and home production. This has recently been highlighted by [Chiappori and Meghir \(2015\)](#) and [Cherchye et al. \(2018\)](#), who propose to compare the economic well-being across individuals in terms of money-metric welfare indices, which take the full basket of goods into account. Thus, in addition to consumption inequality, I also consider inequality the utility ($u(c_{it}, C_{it}, l_{it}, D_{it})$). Note that this requires a structural approach, since utility depends on the calibrated preference parameters, which attach weights to each good and are needed to be able to summarize all goods in a single index. While the main analysis focuses on inequality in utility di-

rectly, I also converted utility levels to a money-metric index, the values of which are more easily interpretable, and obtained similar results for inequality. These robustness checks are relegated to section 5.5.

Finally, I also consider inequality in expected *life-time* utility. The motivation for this is that this is the broadest utility index, which does not only capture the current economic well-being, that may be influenced by short-term shocks, but also the expectation about the future. The expected *remaining* life-time utility of an individual with age t is denoted as $\mathbb{E}u_{i,t}$:

$$\mathbb{E}u_{it} = u(c_{it}, C_{it}, l_{it}, D_{it}) + \beta V_{it}$$

Here, V_{it} denotes the continuation value, conditional on being single or married and the corresponding state variables. To focus on economic inequality, I subtract the expected utility from 'love' (θ) from the continuation value, so that $\mathbb{E}u_{i,t}$ only captures expected streams of consumption, leisure and home production.³³ Note that differences in (remaining) life-time utility are partly due to a difference in the remaining number of periods when individuals do not have the same age (for example, consider comparing $\mathbb{E}u_{it}$ of a 20-year-old and a 40-year-old). To make utility levels comparable across age groups, I compute the *per-period* value u_{it}^{LT} that would be required over the remaining life-time to achieve utility $\mathbb{E}u_{it}$:

$$\sum_{t'=t}^T \beta^{t'-t} u_{it}^{LT} = \mathbb{E}u_{it}$$

Thus, u_{it}^{LT} converts the expected life-time utility to a per-period average and can be interpreted as a measure of expected economic well-being.

5.2 Decomposition Of Inequality

The variance of these outcomes, denoted as X_i , can be decomposed into components due to inequality within and between households and due to singles.³⁴ The variances within and between couples are defined as follows (h is the index of the married household):

$$\begin{aligned} V^W(X_i) &= E\left(V(X_i \mid i \in h)\right) \\ V^B(X_i) &= V\left(E(X_i \mid i \in h)\right) \end{aligned}$$

³³Since θ enters utility additively, one can compute the expected value of future realizations of θ for each state and subtract this value from the continuation values of singles and married individuals as defined previously. Note that the match quality shock θ would in principle introduce some interesting considerations about inequality. Low-wage individuals might be more willing to accept and remain in 'bad' marriages for economic reasons, so that their poverty would be reflected in the non-economic match quality component.

³⁴In appendix B, I also consider other inequality measures, such as the Theil index, for which the results are very similar to the results from using the variance.

The interpretation of the within-couple variance is that it measures how much allocations differ between spouses and takes the expectation of this variance over all couples. The between-household component is the variance of the household means, measuring inequality across households.

To decompose inequality of the population, one further needs to take singles into account. The shares of singles and married individuals in the population are p_s and p_m . Singles enter through two separate terms. The first is the variance of X_i in the group of singles, $V^S(X_i)$. The second term reflects that singles and married individuals differ in their mean outcome.³⁵ This is denoted as the variance ‘between married and single individuals’:

$$\begin{aligned} V^{BMS}(X_i) &= p_s \left(E(X_i | \text{singles}) - E(X_i) \right)^2 \\ &+ p_m \left(E(X_i | \text{married}) - E(X_i) \right)^2 \end{aligned}$$

V^{BMS} partly captures economies of scale. Married individuals can share the public consumption good with their spouse and are able to afford a higher living standard than singles. In addition, spouses can pool home production time and enjoy more leisure. Finally, the term can also reflect selection into marriage - if, for example, high ability individuals were more likely to get married, this would lead to a mean difference between singles and married individuals even when there are no economies of scale.

Using these expression leads to the following decomposition of the population variance:

$$V(X_i) = p_m V^W(X_i) + p_m V^B(X_i) + p_s V^S(X_i) + V^{BMS}(X_i)$$

Note that the components are weighted by p_s and p_m . The higher the fraction of singles or married, the higher is the contribution of the group to the population variance.

Table 4 shows the results of the decomposition based on the calibrated model. The decomposition captures cross-sectional inequality, since it pools all age groups and reflects the population in steady state. The first column reports the decomposition of log private consumption. The single component accounts for 21.4% of the total variance. Note that this value takes the population fraction of singles into account, which reduces the size of the component relative to the variance of singles. The within household variance accounts for 23.9% of inequality in private consumption, while the between household component accounts for about 54.3%. The variance between married couples and singles is small (0.4%), indicating that the mean difference in private consumption between the two groups is small.

The second column shows the decomposition for the per-period utility from private and public consumption, leisure and home production. Thus, this measure takes all private and public goods

³⁵This term is the between-variance that results from applying a within-between decomposition to the two groups of singles and married individuals:

$$V(X_i) = p_m V^M(X_i) + p_s V^S(X_i) + V^{BMS}(X_i)$$

Here, $V^M(X_i)$ is the variance across married individuals.

TABLE 4: Variance decomposition

	$\log(c_{it})$	$u(c_{it}, C_{it}, l_{it}, D_{it})$	u_{it}^{LT}
Variance	0.1	1.26	0.64
Within couples (% of total)	23.9	6.1	9.5
Between couples (% of total)	54.3	65	72.9
Singles (% of total)	21.4	20	17.2
Betw. sin. and mar (% of total)	0.4	8.9	0.4

Notes: The table shows the variance decomposition for private consumption, the per-period utility and (per-period) remaining life-time utility based on the model. Each row reports the fraction of the total fraction due to this component.

into account and weights them according to the calibrated preference parameters. Reflecting the importance of public goods (public consumption and home production), the within-couple component now accounts for 6.1% of total inequality. Interestingly, the variance between singles and married couples gets more important in this case (8.9%), which shows that the economies of scale in terms of public consumption and home production, which couples enjoy relative to singles, are substantial.

The third column finally shows the decomposition for expected life-time utility (u_{it}^{LT}). At any point in time, some individuals may be worse (or better) off due to randomness in wages, marital status and fertility, but this is mitigated when taking the expectation over the future. Note that the presence of divorce creates additional differences in the well-being of spouses. For example, consider a couple of a low-wage man and a high-wage woman, which is likely to divorce soon. If marriages on average are substantially assortative, the man would be worse off than the woman in terms of expected future utility. This reasoning also applies to inequality *across* couples, since a 'stable' couple is better off than a couple where divorce is impending. Overall, the within-household components accounts for 9.5% of total inequality for the life-time utility measure, which is relatively similar to the case of looking at per-period utility only. Note that the importance of inequality between singles and married individuals drops significantly (from 8.9% to 0.4%). This reflects singlehood being a transitory phenomenon for most individuals. While singles have a lower living standard than married individuals, they mostly expect to get married soon.

For comparison and robustness, I also conducted the decompositions for alternative measures, including measures that convert utility levels into money-metric indices. These cases are discussed in more detail in section 5.5. The results from the alternative measures are very similar to considering inequality in utility.

5.3 Experiment: Increase in Tax Progressivity

The policy experiment studies the effects of a hypothetical increase in progressivity, by varying the progressivity parameter (ψ_2) of the tax function

$$T(y) = \max \left((1 - \psi_1 y^{\psi_2}) \cdot y, 0 \right)$$

The level parameter ψ_1 is adjusted to keep the budget of the government balanced. The policy experiment thus raises tax rates for higher incomes, while lowering them for lower income. The current system is approximated by a progressivity parameter of 0.15, which results from fitting the tax function to income tax rates.³⁶ When the progressivity parameter (ψ_2) is changed, the level parameter (ψ_1) is adjusted to ensure that the budget constraint of the government is balanced. The target level of expenditure is set to the revenue that the government obtains for the current tax system in the calibrated model. To study the effects of these reforms, I compare steady states, which can be interpreted as analyzing the long-run impact of the reform.³⁷

The hypothetical reform increases the progressivity parameter by 0.06. To illustrate the magnitude of the tax change, figure 4 shows the change of the average tax rate due to the policy change. The reform reduces the average tax rate at lower incomes by up to 4%, while it increases the average tax rate for high income levels (around 150.000 €) by $\approx 5\%$. The motivation for this particular reform is that it is within the range of the cross-country variation in progressivity estimated e.g. by [Holter, Krueger, and Stepanchuk \(2019\)](#). The focus of the policy analysis is mainly to compare the relative magnitude of the different components of inequality. I also conducted the analysis for different potential reforms and the relative importance of the different components (in particular, the role of intra-household inequality) is similar.

5.4 How Does The Reform Affect Inequality?

To analyze how the reform affects inequality along the dimensions of the marriage market, I decompose the *change* in inequality. The variance before and after the reform are denoted as V_0 and V_1 :

$$V_k(X_i) = p_{m,k} V_k^W(X_i) + p_{m,k} V_k^B(X_i) + p_{s,k} V_k^S(X_i) + V_k^{BMS}(X_i)$$

The goal is to express the percentage change in the variance ($\hat{V} = \frac{V_1 - V_0}{V_0}$) as a weighted sum of the change in each of the variances from the decomposition. The *change* in the variance between the

³⁶I fit income tax rates, rather than the entire tax and transfer system, since transfers are typically means-tested on the family level, whereas the income tax is assessed individually.

³⁷Computing the transition path would raise some complications regarding the marriage market equilibrium. Following the reform, individuals would have to forecast how the reform affects the marriage and divorce decisions of others, and thereby the pool of singles they will meet. Intuitively, the main difference between the transition and the new steady state concerns the decision weight in existing couples: changes in the weight would require that a participation constraint binds, which is more likely for larger reforms.

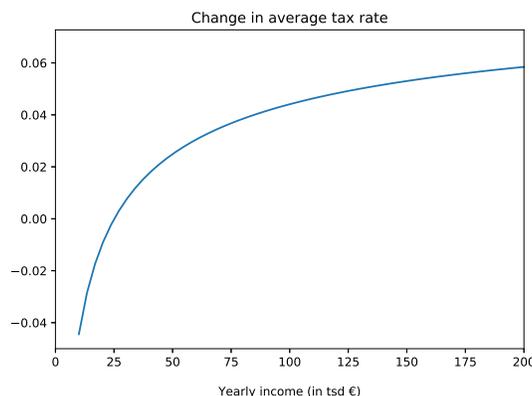


FIGURE 4

Notes: The figure shows the change in the average tax rate due to the reform.

two policy regimes can be decomposed in the following way:

$$\hat{V}(X_i) = \omega_1 \hat{V}^S + \omega_2 \hat{V}^W + \omega_3 \hat{V}^B + \hat{C}^{BMS} + \omega_0 \hat{p}_m + \hat{R}$$

\hat{V}^S is the percentage change in the variance of singles, \hat{V}^B and \hat{V}^W are the changes of the variances between and within married couples. These changes are weighted by the importance of the component before the reform ($\omega_1, \omega_2, \omega_3$), to reflect that each of the percentage changes acts on a different base. For example, consider a case where the within-couple component declines by 20% and accounts for 10% of the total variance before the reform. This, in isolation, would lead to a 2% ($= 0.1 \cdot 0.2$) decrease in the total variance.³⁸ I further separate the effect of changes in the fraction of the population that is married ($\omega_0 \hat{p}_m$). p_m enters all of the components of the decomposition and has two effects. First, an increase in p_m reduces the weight placed on the variance across singles and increases it for the variances within and across couples. Second, p_m enters V^{BMS} . Separating the effect of p_m introduces a residual into the decomposition.³⁹ Finally, \hat{C}^{BMS} is based on the variance between couples and singles. The expressions for this term and each of the weights are described in more detail in appendix C.

Table 5 shows the results of the decomposition for private consumption and the utility measures. The first row contains the reduction in inequality due to the reform (for example, 10.9% for private consumption). The other rows decompose the total reduction into parts due to each of the components. For private consumption, 24.77% of the total reduction is due to a reduction of within-household inequality. Thus, a fairly substantial part of the total reduction in private consumption would be missed by assuming equal sharing within households. For the utility measures, the corresponding fractions are 11.43% and 9.84%, which is smaller than for private

³⁸Formally, this corresponds to the formula $x \hat{+} y = \hat{x} \frac{x}{x+y} + \hat{y} \frac{y}{x+y}$.

³⁹The residual results from applying the approximation $\hat{x}\hat{y} = \hat{x} + \hat{y} + \hat{x}\hat{y} \approx \hat{x} + \hat{y}$.

consumption only, but still a noteworthy contribution to the total reduction in inequality. The role of the between-couple component is largest and ranges between 54.29% and 70.49%. The decomposition further allows to quantify the role of changes in the fraction of singles. The fraction of singles shifts the weights of the variance components in the decomposition. This, in itself, can generate a change in inequality. For example, when the variance of singles is higher than the variance of married individuals, the lower fraction of singles can decrease inequality. Based on the calibrated parameter values, the fraction of singles declines by a very modest amount (0.2 percentage points) in the more progressive tax system, which leads to very small implications for inequality (up to 2.86% of the total reduction).

TABLE 5: Variance decomposition - Policy change

	$\log(c_i)$	$u(c_i, C_i, l_i, D_i)$	u_{it}^{LT}
Change (%)	-10.9	-3.5	-6.1
Decomposition of change			
Within couples (%)	24.77	11.43	9.84
Between couples (%)	63.3	54.29	70.49
Singles (%)	11.93	31.43	21.31
Betw. sin. and mar. (%)	-0.92	0	0
Single probability (%)	0	2.86	-1.64

Notes: This table shows the effect of the policy change on the variance of private consumption, the per-period utility and life-time utility. The first row shows the total change. The other rows decompose this total change in inequality and add up to 100%.

The decompositions from tables 5 and focused on the contributions of each component to the total change in variance. This depends both on how important a component is under the status quo and how *reactive* it is to the tax system. Recall the example where the within-couple component declines by 20% (this is the 'reactiveness') due to the reform and accounts for 10% of inequality under the status quo, resulting in a 2% decrease in inequality. Table 6 focuses only the *reactiveness*, in terms of how strongly the within- and between-couple variance decline. This is interesting because it could be that the reactivity of the components differs, which could be interpreted as progressivity being more 'effective' at reducing either the within- or between couple variance.

Table 6 shows that the reactivity of both components is quite similar for private consumption. The between-couple component declines a bit more strongly than the within-couple com-

ponent (-12.7% vs -11.4%). Thus, the relative rate of decline is 0.89. Interestingly, the within-couple variance declines *more strongly* than the between-couple component for the per-period utility (-6.1% vs -2.7%), so that in this case the within-couple component is more reactive. Note that differences in reactivity imply that the *composition* of inequality is endogenous to the tax system. For example, for the utility measure, the fact that the within-household component declines more strongly than the between-household component suggests that a smaller fraction of the *remaining* inequality should be due to within-household inequality. In practice, however, these composition changes are very small, which is shown in appendix table A5. The share of the within-household component declines only from 6.2 to 6.1%. Finally, table A6 replicates the decomposition of the effect of increasing progressivity for the case of *decreasing* it to a level of $\psi_2 = 0$, which replaces the tax function by a proportional tax. The relative importance of the changes in within- and across-couple inequality are remarkably similar to the previous results. Introducing a flat tax increases inequality and within-household inequality accounts for 25.4% of this increase for private consumption and 8.33% of the increase for the per-period utility.

TABLE 6: Unweighted change in variance

	Var. within couples	Var. between couples	Within rel. to between
(a) $\log(c_i)$			
Variance ($\psi_2 = 0.15$)	0.029	0.066	0.439
Change ($\psi_2 = 0.21$)	-0.114	-0.127	0.898
(b) $u(c_i, C_i, I_i, D_i)$			
Variance ($\psi_2 = 0.15$)	0.095	1.033	0.092
Change ($\psi_2 = 0.21$)	-0.061	-0.027	2.259

Notes: This table shows the variance within and between couples and of singles. The first row in each section shows the level of the variance for the calibrated model and the other rows show the change (in percent) due to policy reforms. The final column shows the ratio between the within and the between component of each row.

5.5 Inequality Decompositions: Robustness

For further comparison and robustness, table 7 reports the results of the variance decomposition for three additional measures of well-being. The first column reports the decomposition results for inequality in total consumption expenditure ($c_{it} + C_{it}$). The first noteworthy feature is

that the fraction due to the within-couple components drops significantly relative to the case of private consumption only and becomes small (1.3%). This reflects household expenditure being to an important extent public and indicates that there is a limited role for intra-household inequality in total consumption, unless one makes further assumptions about preference heterogeneity or reduces the share of consumption that is considered public. The importance of the within-household component here is lower than found in [Lise and Seitz \(2011\)](#), partly reflecting that they assume a lower fraction of public consumption. The second interesting feature is that the variance between singles and married individuals becomes much more important for total expenditure than for private expenditure only, which is due to economies of scale in consumption.

The motivation for the remaining three columns is that the main analysis has focused on inequality in *utility*, the cardinal value of which is hard to interpret. Thus, one might wonder if the results would differ for money-metric measures of well-being. The second column computes the variance decomposition for 'full consumption', as in [Lise and Yamada \(2018\)](#). Full consumption is defined as the sum of consumption expenditure and the *market value* of leisure and home production time. This is defined as follows:

$$I_i = c_i + C_i + \tilde{w}_i l_i + \tilde{w}_i d_i + \tilde{w}_j d_j$$

Here, \tilde{w}_i is the net hourly full-time wage of the individuals, which captures the opportunity cost of leisure and home hours, and the index j refers to the home hours of the partner. Note that tax reforms have a direct impact on net wages and therefore influence this measure through their impact on \tilde{w}_i and \tilde{w}_j in addition to the impact on allocations. Still, the measure can be considered as a useful reference and comparison point, especially since it can in principle be measured in the data more easily than measures which require estimates of preferences.

The third and fourth column is an equivalence scale in the spirit of [Chiappori and Meghir \(2015\)](#) or [Pendakur \(2018\)](#). These measures are based on comparing utility levels across people, but convert the utility level of each individual into the monetary amount that a reference person would need to achieve this utility level. For the utility level of each person ($u(c_i, C_i, l_i, D_i)$), one can compute the equivalent amount of resources the reference person would need in order to obtain this utility level. The reference person is set to be a childless man, who does not work and receives the equivalent amount as a transfer, while optimally choosing leisure and home time. The fourth column applies this approach to the expected life-time utility and computes the amount that the reference person would need to obtain the expected per-period utility u_{it}^{LT} .

The overall take-away from considering these measures in [table 7](#) is that the results are very similar to considering inequality in utility ($u(c_{it}, C_{it}, l_{it}, D_{it})$ and u_{it}^{LT}) directly. In particular, this holds for the role of within-household inequality. Recall that for the per-period utility, 6.1% of inequality is due to the within-component. For the equivalence scale and full expenditure, this value is 6.7% and 5.6%. Similarly, the equivalence scale for life-time utility leads to a very similar result as previously. [Table 8](#) further shows the decomposition of the reduction in inequality due to these measures, which are also very similar to the previous decompositions. For example, for

the per-period utility, 11.43% of the total reduction is due to the within component, while the corresponding number is 11.1% for the equivalence scale.

TABLE 7: Variance decomposition - alternative measures

	$\log(c_i + C_i)$	Log full exp.	Log eq. sc.	Log eq. sc. (LT)
Singles (%)	20.4	22.6	14.2	17.3
Within couples (%)	1.3	5.6	6.7	9.0
Between couples (%)	52.0	59.5	70.2	73.5
Betw. sin. and mar (%)	26.3	12.4	9.0	0.3

Notes: This table computes the variance decomposition for the status quo for the alternative measures and reproduces table 4 for these measures. The columns contain log total consumption, log full expenditure and the logs of the equivalence scale values corresponding to $u(c_{it}, C_{it}, l_{it}, D_{it})$ and u_{it}^{LT} .

TABLE 8: Policy change - alternative measures

	$\log(\mathbf{c}_i + \mathbf{C}_i)$	Full expenditure	Log eq. sc.	Log eq. sc. (LT)
Total (%)	-8.4	-10.9	-4.5	-7.7
Decomposition of change				
Within couples (%)	1.2	4.6	11.1	10.4
Between couples (%)	81	60.6	62.2	68.8
Singles (%)	15.5	28.4	24.4	20.8
Betw. sin. and mar. (%)	-1.2	4.6	0	0
Single probability (%)	3.6	1.8	0	-1.3

Notes: This table shows the effect of the policy change for different outcomes. . The columns contain log total consumption, log full expenditure and the logs of the equivalence scale values corresponding to $u(c_{it}, C_{it}, l_{it}, D_{it})$ and u_{it}^{LT} . The first row shows the total change. The sum of the components from other rows is equal to the total change in the first row.

A final question for robustness is whether the choice of the inequality measure matters, which so far was the variance. Similar decompositions can be applied to the Mean Logarithmic Deviation and the Theil Index, since they are subgroup-decomposable. The results are qualitatively similar and relegated to appendix B.

5.6 Marriage outcomes

In this section, I discuss marriage outcomes and how they are affected by the increase in tax progressivity in more detail, focusing on differences by ability and gender asymmetries. To facilitate the interpretation, the analysis groups individuals into ‘ability groups’ ranging from low (L) to high (H). Each of these groups contains averages over three types of the actual type distribution.⁴⁰ Table 9 shows marriage rates and the mean Pareto weight for women and men from each of these groups. The table is based on considering all meetings between individuals below age 30, which is the period in which most marriages occur, and calculating the probability that the meeting results in a marriage, and the mean Pareto weight conditional on marriage. For men, the table reports *their* Pareto weight (i.e. $1 - \lambda_{it}$).

TABLE 9: Marriage rates and Pareto weights

	Women - L	W - LM	W - UM	W - H	Men - L	M - LM	M - UM	M - H
(a) Status quo								
Marriage probability	0.289	0.274	0.254	0.223	0.227	0.251	0.271	0.289
Own Pareto weight	0.453	0.472	0.496	0.526	0.451	0.487	0.528	0.584
(b) Reform								
Change in mar. prob. (in pp)	-0.22	-0.26	0.29*	0.7***	0.38**	0.43**	-0.29*	0.05
Change in PW (in pp)	0.31***	0.45***	0.52***	0.62***	-0.11*	-0.38***	-0.43***	-0.94***

Notes: The table reports the marriage market outcomes under the status quo and after the reform. The asterisks indicate whether the difference of each outcome between the two simulations is statistically significant at the 1%, 5% or 10% level.

The table shows an interesting gender asymmetry with respect to marriage rates. Low ability women are *more* likely to get married than high-ability women, whereas low-ability men are *less* likely to get married than high ability men. Quantitatively, these effects are modest and the marriage probability varies by up to 6 percentage points across the ability groups. The reason for this effect is that marriages in which men are from a higher ability group than the women have a higher surplus than the opposite combination, since women are more likely to reduce their market hours in marriage. For both men and women, the Pareto weight is increasing in the own ability type and the average ranges between 0.45 and 0.58.

The second part of the table shows how these outcomes differ in the new steady state under the more progressive tax system. The changes are expressed in percentage points. Since some of these are quite small, the table also reports whether the difference is statistically significant from zero, rather than representing simulation noise. The changes in the Pareto weight are mostly significant on the 1% level. Women on average receive higher Pareto weights, whereas the opposite is the

⁴⁰For example, group L contains individuals of type 1, 2 and 3.

case for men.⁴¹ Regarding marriage rates, only the rate of the highest two groups of women and the lowest two groups of men increase in a statistically significant way. The magnitude of these changes ranges between 0.29 and 0.7 percentage points. To further illustrate these patterns, table 10 shows how the marriage probability and the Pareto weight change for each *combination* of types - for example when a women from the highest group meets a men from the lowest group. While the magnitude of these changes is again modest, the table reveals some further gender asymmetries. For marriage rates, most of the significant changes occur for meetings between high ability women and low ability men, whereas there is no significant change for the high ability men and low ability women. For example, the reform increases the probability that a women from the highest group gets married to a men from the lowest group by 0.64 percentage points. For Pareto weights, the pattern is different since the effects are largest for low ability women getting married to high ability men. The Pareto weight of the woman increases by up to 1.52 percentage points (recall that Pareto weights are scaled between zero and one).

TABLE 10: Changes in marriage rates and Pareto weights - by type combination

	Men - L	Men - LM	Men - UM	Men - H
(a) Marriage probability				
Women - L	-0.02	-0.04	0.08	-0.22
Women - LM	0.2	-0.87**	-0.21	-0.13
Women - UM	0.46**	0.78***	-0.34	0.17
Women - H	0.64***	0.83***	0.41*	0.59
(b) Pareto weight				
Women - L	-0.06	0.74***	0.96***	1.52***
Women - LM	0.28	0.17***	0.91***	1.4***
Women - UM	-0.24	0.45***	0.31***	1.45***
Women - H	-0.39*	-0.62***	0.07	0.59***

Notes: The table reports changes between the status quo and the new steady state, expressed in percentage points. In this table, the Pareto weight always refers to the weight of the woman. The asterisks indicate whether the difference of each outcome between the two simulations is statistically significant at the 1%, 5% or 10% level.

⁴¹This results from the gender wage gap, since each type of woman on average have a lower wage than the corresponding type of man.

5.7 Relating Changes in Marriage and Divorce to Inequality

Changes in marital behavior can contribute to the effect of the reform on inequality. For example, a decrease in the assortativeness of matching would decrease inequality *across* couples, since more high-income individuals get married to low-income individuals. At the same time, this would increase inequality *within* couples, since the fraction of couples with unequal earnings potential would rise. The goal of this section is to quantify the role of such composition effects and analyze whether they contribute to the results from the variance decompositions of section 5.4.

I perform further decomposition of each of the components of the variance decomposition, focusing on the role of the ability types. To illustrate the general approach, I first focus on the variance across couples. The share of couples, in which the wife has type k and the husband has type j is denoted as s_{kj} . The across-couple variance can be rewritten in the following way:

$$\begin{aligned} V(\bar{X}_h) &= \sum_{kj} s_{kj} V(\bar{X}_h|k, j) + \sum_{kj} s_{kj} \left(E(\bar{X}_h|k, j) - E(\bar{X}_h) \right)^2 \\ E(\bar{X}_h) &= \sum_{kj} s_{kj} E(\bar{X}_h|k, j) \end{aligned}$$

Here, \bar{X}_h is the mean outcome for couple h (recall that the across-couple variance is the variance of the household means). Thus, the variance across couples can be expressed as a function of the variance among couples of type (k, j) ($V(\bar{X}_h|k, j)$), the mean outcome of couples of this type ($E(\bar{X}_h|k, j)$) and the type share s_{kj} . The idea of the decomposition is to use this formula to compute a counterfactual post-reform variance, in which only the type shares s_{kj} are varied to their post-reform level and all the other components (variance and mean conditional on type) remain at their pre-reform values. The interpretation of this procedure is to ask to what extent the change in the type probabilities alone can explain the total change in the variance.⁴²

Similar formulas can be derived for the other variance components. The variance across singles can be analyzed completely analogous to the variance across couples, using the same formula and only replacing the household mean (\bar{X}_h) by the value of the single (X_i). The corresponding formula for the within-couple variance is:

$$E\left(V(X_i|i \in h)\right) = \sum_{i,j} s_{k,j} E\left(V(X_i|i \in h) | k, j\right)$$

Like for the case of the across-couple variance, changes in the within-couple variance are in principle driven by the type shares ($s_{k,j}$) and the variances conditional on the types. Finally, the variance between singles and married individuals can be further decomposed using the following formu-

⁴²An alternative approach which would yield very similar results would be to reweight the simulated post-reform data to have the same composition of observables (in terms of types k and j) as the pre-reform data and compute the variance decomposition based on the reweighted data, as is done in [Lise and Seitz \(2011\)](#).

las:

$$\begin{aligned}
V^{BMS}(X_i) &= (1 - p_m) \left(E(X_i|S) - E(X_i) \right)^2 + p_m \left(E(X_i|M) - E(X_i) \right)^2 \\
E(X_i|S) &= \sum_k s_k^S E(X_i|S, k) \\
E(X_i|M) &= \sum_{kj} s_{k,j}^M E(X_i|M, k, j) \\
E(X_i) &= (1 - p_m) E(X_i|S) + p_m E(X_i|M)
\end{aligned}$$

s_k^S is the fraction of singles of type k and $s_{k,j}^M$ is the fraction of couples with types k and j . The common feature of these formulas, for each of the components of the variance decomposition, is that they allow to vary only the type composition of the groups of married and single individuals.⁴³

Table 11 shows the results for this decomposition for inequality in private consumption and the per-period utility. The interpretation of the first row ('fixed p_m ') is that it adjusts s_{jk} to the values in the new steady state, but keeps everything else (including p_m) constant. The interpretation is that only the type composition of the pools of married and single individuals are varied. The second row also lets p_m adjust and considers changes in the *weighted* contribution to the total variance, i.e. each of the following four summands:

$$V(X_i) = p_m V^W(X_i) + p_m V^B(X_i) + p_s V^S(X_i) + V^{BMS}(X_i)$$

For example, the interpretation of the very first table entry is that changes in s_{jk} explain -1.91% of the reduction in the within-couple variance, which means that these changes by themselves would marginally increase within-couple inequality. The overall take-away from considering the 'fixed p_m ' case for both outcome variables is that the effect of composition effects alone is very small. Only for the variance between married and single individuals (BMS), the composition effect can generate a change of a noteworthy magnitude. An interesting aspect of these results is that the results from the previous section suggested that the increase in tax progressivity could decrease assortative mating, since there was a small increase in the likelihood of low-ability men getting married to higher-ability women. However, this effect is not strong enough to lead to a quantitatively remarkable reduction in inequality.

⁴³Note that this is an accounting decomposition, since endogenous quantities are varied independently.

TABLE 11: Fraction of variance change explained by type probabilities (in %)

	Within couples	Between couples	Singles	BMS
(a) $\log(c_{it})$				
Fixed p_m	-1.91	-0.37	2.75	24.45
Flexible p_m	-3.53	-1.8	20.01	13.99
(b) $u(c_{it}, C_{it}, I_{it}, D_{it})$				
Fixed p_m	-2.8	-0.29	1.34	-23.95
Flexible p_m	-5.82	-7.79	14.91	37.75

Notes: The table reports the share of the change in each of the variance components that can be explained only by the change in type probabilities.

When also letting the fraction of married individuals adjust ('flexible p_m '), the explained fraction become somewhat larger. Note that the outcome is now for example $p_m V^W(X_i)$ (instead of $V^W(X_i)$). Since the reform increases the fraction of married individuals, this in itself increases the contribution of inequality within and across couples and decreases the contribution of singles. Thus, the explained fraction is negative in the first two columns and positive in the third. The explained share is overall largest for the per-period utility and ranges between -7.8% and 14.91% .

5.8 Social welfare

In this section, I show how the reform affects expected life-time welfare. While inequality is strongly linked to social welfare, the most natural social welfare function in my model is based on the expected life-time welfare of a newly born cohort. Thus, the analysis in this section is concerned with the overall welfare impact of the reform, and to what extent it is driven by marriage market adjustments.

5.8.1 The comparison case

To assess the role of the marriage market, I consider an alternative version of the model in which the decision rules for marriage, divorce and the Pareto weight are exogenous and set to exactly reproduce the decision rules of the calibrated model. Then, I simulate the reform in this alternative model. The interpretation is to ask what the effects of the reform would be if marriage and divorce decisions and the Pareto weights were exogenous and would not react to the change

in tax policy. I will refer to the initial equilibrium given this policy parameter as the benchmark case/schedule. In addition, I will refer to the model with fixed decisions as the *alternative* or *fixed marriage market* model.

More concretely, suppose that the initial progressivity parameter is ψ_2^0 . Solving the model leads to decision functions which map the state variables of two singles that meet, ω_f and ω_m , and the realization of the match quality shock θ into a decision whether to get married ($M \in \{0, 1\}$) and an initial Pareto weight $\lambda \in [0, 1]$:

$$\begin{aligned} M &= m_0(\omega_f, \omega_m, \theta) \\ \lambda &= m_1(\omega_f, \omega_m, \theta) \\ \omega_f &= (\epsilon_f, \alpha_f, b_f, A_f) \\ \omega_m &= (\epsilon_m, \alpha_m, b_m, A_m) \end{aligned}$$

The state variables ω_g comprise the productivity shock, the type, the presence of children, and assets. Similarly, there are decision functions d_0 and d_1 which map the state variables of any couple (ω_c) into the decision whether to get divorced ($D \in \{0, 1\}$) and whether to renegotiate the current Pareto weight λ to a new Pareto weight $\tilde{\lambda}$:

$$\begin{aligned} D &= d_0(\omega_c) \\ \tilde{\lambda} &= d_1(\omega_c) \\ \omega_c &= (\theta, \epsilon_f, \epsilon_m, \alpha_f, \alpha_m, b, A, \lambda) \end{aligned}$$

In the comparison case, I solve the model for a new tax schedule ψ_2^1 , but use the decision functions m_0, m_1, d_0, d_1 that were obtained given ψ_2^0 , instead of allowing individuals to determine these outcomes endogenously. In particular, this procedure does not allow the Pareto weights to adjust following a reform. Instead, spouses are exogenously assigned the 'old' Pareto weight, which would have been the bargaining outcome given the old tax schedule. As a result, comparing the alternative and the full model allows to study the implications of the endogenous adjustment of the Pareto weights. The exercise is related to the exercise conducted in Knowles (2012), who compares a model with flexible and fixed Pareto weights in order to quantify the error one would make with a unitary version of the model relative to a model with endogenous Pareto weights. Note that in the alternative model used here, the exogenous Pareto weight depends on the characteristics of the couple and is allowed to change exogenously over time.⁴⁴

The alternative model has different implications for how intra-household inequality changes

⁴⁴Further note that when there are no dynamic choices (savings), the alternative model always delivers exactly the same marriage market outcomes that are obtained under the benchmark schedule. In other words, the cross-sectional distribution of marriages and Pareto weights is *identical* to the benchmark model. With dynamic choices, policy changes can still affect marriage and divorce rates and Pareto weights even when decision functions are fixed. This is the case because decisions are fixed *conditional on state variables* and state variables (i.e. assets) are partly endogenous and affected by the reform. Thus, the interpretation of the exercise is how much the endogenous *adjustment of the decision functions* matters for welfare, rather than marriage market outcomes per se.

with progressivity. When Pareto weights are fixed, the tax system can influence intra-household inequality mainly by changing the relative leisure of spouses. A less progressive system can increase leisure of secondary earners with low Pareto weights if it leads to a reduction in their labor supply and home production hours increase less than work hours are increased. With flexible Pareto weights, progressivity directly affects relative consumption by changing the Pareto weight.

5.8.2 Results

Table 12 shows the overall welfare change for the increase in progressivity for each ability type, averaged over men and women. These are also illustrated graphically in figure 5. The welfare change is expressed in terms of equivalent variation. This is the percentage change in consumption in each state under the status quo, which makes individuals indifferent between the status quo and the reform.⁴⁵ The second and third columns show the equivalent variation (in percent) for the full model and for the comparison model and the fourth column contains the difference between the two. This difference can be interpreted as the effect of the endogenous marriage market adjustments: it is the equivalent variation that the full model generates relative to the case of fixed marriage market decisions. To ease interpretation of these numbers, the final column converts the difference to a percentage amount, relative to the welfare impact in the fixed marriage market case. Note that "1" is the lowest ability type and "12" is the highest.

As one would expect, the lowest ability types benefit from the increase in progressivity, whereas the highest ones lose, both for the full model and the fixed marriage market. In the middle of the distribution most types are made worse off. This can be the case because of the labor supply disincentives generated by a more progressive tax system. Note that the fact that most types lose is also reflected in the aggregate welfare change, which averages over all types and therefore corresponds to a utilitarian welfare objective. The aggregate welfare loss is -0.92% for the full model and -1.22% for the fixed marriage market model. Thus, in relative terms, the difference amounts to 24% of the impact without marriage market adjustments.

The welfare difference between the full model and the fixed marriage market is largest for the two lowest types. In particular, the lowest type experience a welfare gain of 1.98% of consumption in the full model, whereas it would only be 0.86% if marriage market decision are fixed. These differences get somewhat smaller for higher ability types. In the upper half of the ability distribution, the marriage market adjustments generate between -12 and 15% of the welfare impact with a fixed marriage market. In general, these values illustrate the importance of the marriage market relative to the 'standard' channels of tax progressivity and show that while the marriage market does not dominate the welfare calculation, it contributes in a significant way.

⁴⁵Since the model features both private and public consumption, the equivalent variation is decomposing by increasing consumption of both of these goods by a certain percentage amount in each state of the world.

TABLE 12

	(1) Full model	(2) Fixed MM	Difference: (1) - (2)	Relative
1	1.98	0.86	1.12	1.31
2	1.37	0.53	0.84	1.61
3	-0.03	-0.29	0.26	-0.89
4	0.37	-0.25	0.62	-2.50
5	0.04	-0.49	0.54	-1.09
6	-0.47	-0.70	0.22	-0.32
7	-0.92	-1.29	0.36	-0.28
8	-1.46	-1.66	0.19	-0.12
9	-1.88	-1.74	-0.15	0.08
10	-2.62	-2.87	0.25	-0.09
11	-3.48	-3.27	-0.21	0.07
12	-3.94	-3.44	-0.51	0.15
Total	-0.92	-1.22	0.30	-0.24

Notes: This table shows the difference in the welfare change between the full model and the fixed marriage market case. It is based on the values shown in figure 5.

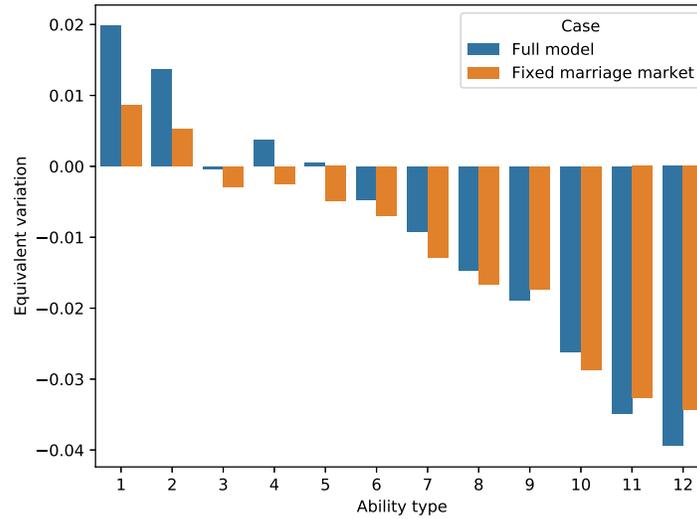


FIGURE 5: Welfare impact along the ability distribution

Notes: The figure shows the change in welfare (expressed as percent of consumption) for the full model and the fixed marriage market case.

6 Conclusion

This paper studies how the progressivity of the labor income tax affects inequality through the marriage market. Progressive taxation reduces intra-household inequality because it makes the relative outside options of spouses - the values of being single - more equal. In addition, single rates, marriage and divorce are also affected since progressivity changes the economic value of individuals on the marriage market and how selective they are about potential partners. To study these effects, I calibrate an equilibrium model of marriage, divorce, labor supply and savings to

data from the Netherlands. The main question that the analysis addresses is to what extent the marriage market channels - bargaining and marriage - affect inequality on top of the usual effects of progressivity.

The model is first used to decompose inequality in consumption and welfare into components due to within and between married couples. The within-couple component captures that spouses consume different amounts of private goods and has received little attention in studies of progressivity. In the calibrated model, intra-household inequality accounts for 23.9% of the cross-sectional variance in private consumption. The model further allows to study inequality in the utility from private and public consumption, leisure and home production. In this case, the intra-household component accounts for 6.1% of the total variance. The model is then used to study a hypothetical reform that increases progressivity by 40% relative to its current level. The contribution of the intra-household component to the total reduction in inequality is 24.77% for private consumption and 11.9% for the utility measure. These quantitative findings shed light on the question of how important intra-household inequality could be relative to overall inequality, which has received little attention in the literature so far. The model further suggests that the induced changes in marriage and divorce have small implications for inequality.

There are several interesting dimensions in which the present analysis could be extended. While my paper has focused on the role of the labor income tax, it would be very interesting to model the transfer system in more detail, which has an important redistributive function at the lower end of the income distribution. Another very interesting extension would be to consider the role of endogenous skill formation. Tax progressivity can reduce the incentive to accumulate human capital in working life and to invest in skills in terms of education and both of these aspects interact with marriage and intra-household bargaining. Finally, while I have focused on intra-household inequality in personal consumption and leisure, allowing for heterogeneous preferences, for example regarding public good expenditure or risk aversion, would increase the scope for disagreement between spouses. This would amplify the importance of intra-household inequality and thereby the welfare implications of changes in bargaining positions.

References

- Alesina, Alberto, Andrea Ichino, and Loukas Karabarbounis. 2011. "Gender-based taxation and the division of family chores." *American Economic Journal: Economic Policy* 3 (2):1–40.
- Attanasio, Orazio P and Valérie Lechene. 2014. "Efficient responses to targeted cash transfers." *Journal of Political Economy* 122 (1):178–222.
- Bastani, Spencer. 2013. "Gender-based and couple-based taxation." *International Tax and Public Finance* 20 (4):653–686.
- Bick, Alexander and Nicola Fuchs-Schündeln. 2017. "Taxation and Labour Supply of Married Couples across Countries: A Macroeconomic Analysis." *The Review of Economic Studies* 85 (3):1543–1576.
- Blundell, Richard, Luigi Pistaferri, and Ian Preston. 2008. "Consumption inequality and partial insurance." *American Economic Review* 98 (5):1887–1921.
- Blundell, Richard and Ian Preston. 1998. "Consumption inequality and income uncertainty." *The Quarterly Journal of Economics* 113 (2):603–640.
- Bronson, Mary Ann and Maurizio Mazzocco. 2019. "Taxation and Household Decisions: an Intertemporal Analysis." *Unpublished manuscript* .
- Chade, Hector and Gustavo Ventura. 2002. "Taxes and Marriage: A Two-Sided Search Analysis." *International Economic Review* 43 (3):955–985.
- . 2005. "Income taxation and marital decisions." *Review of Economic Dynamics* 8 (3):565–599.
- 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, and Frederic Vermeulen. 2012. "Married with children: A collective labor supply model with detailed time use and intrahousehold expenditure information." *American Economic Review* 102 (7):3377–3405.
- Cherchye, Laurens, Thomas Demuynck, Bram De Rock, and Frederic Vermeulen. 2017. "Household consumption when the marriage is stable." *American Economic Review* 107 (6):1507–34.
- Chiappori, Pierre-André, Monica Costa Dias, and Costas Meghir. 2018. "The Marriage Market, Labor Supply, and Education Choice." *Journal of Political Economy* 126 (S1):S26–S72.

- Chiappori, Pierre-Andre and Maurizio Mazzocco. 2017. "Static and intertemporal household decisions." *Journal of Economic Literature* 55 (3):985–1045.
- Chiappori, Pierre-André and Costas Meghir. 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* .
- Conesa, Juan Carlos and Dirk Krueger. 2006. "On the optimal progressivity of the income tax code." *Journal of Monetary Economics* 53 (7):1425–1450.
- de Boer, Henk-Wim, Egbert Jongen, Patrick Koot et al. 2018. "Optimal Taxation of Secondary Earners in the Netherlands: Has Equity Lost Ground?" Tech. rep., CPB Netherlands Bureau for Economic Policy Analysis.
- Fernandez, Raquel, Nezih Guner, and John Knowles. 2005. "Love and money: A theoretical and empirical analysis of household sorting and inequality." *The Quarterly Journal of Economics* 120 (1):273–344.
- Fernández, Raquel and Richard Rogerson. 2001. "Sorting and long-run inequality." *The Quarterly Journal of Economics* 116 (4):1305–1341.
- Fernandez, Raquel and Joyce Cheng Wong. 2017. "Free to leave? A welfare analysis of divorce regimes." *American Economic Journal: Macroeconomics* 9 (3):72–115.
- Galichon, Alfred, Scott Duke Kominers, and Simon Weber. 2018. "Costly concessions: An empirical framework for matching with imperfectly transferable utility." *Journal of Political Economy* (Forthcoming) .
- 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.
- Greenwood, Jeremy, Nezih Guner, Georgi Kocharkov, and Cezar Santos. 2016. "Technology and the changing family: A unified model of marriage, divorce, educational attainment, and married female labor-force participation." *American Economic Journal: Macroeconomics* 8 (1):1–41.
- Guner, Nezih, Remzi Kaygusuz, and Gustavo Ventura. 2011. "Taxation and household labour supply." *The Review of Economic Studies* 79 (3):1113–1149.
- . 2012. "Taxing women: A macroeconomic analysis." *Journal of Monetary Economics* 59 (1):111–128.

- Guvenen, Fatih and Michelle Rendall. 2015. "Women's emancipation through education: A macroeconomic analysis." *Review of Economic Dynamics* 18 (4):931–956.
- Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante. 2017. "Optimal tax progressivity: An analytical framework." *The Quarterly Journal of Economics* 132 (4):1693–1754.
- Holter, Hans A, Dirk Krueger, and Serhiy Stepanchuk. 2019. "How does tax progressivity and household heterogeneity affect Laffer curves?" *Quantitative Economics* .
- Kalai, Ehud. 1977. "Proportional solutions to bargaining situations: interpersonal utility comparisons." *Econometrica: Journal of the Econometric Society* :1623–1630.
- 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.
- 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. 2018. "Household Sharing and Commitment: Evidence from Panel Data on Individual Expenditures and Time Use." *The Review of Economic Studies* .
- Low, Hamish, Costas Meghir, Luigi Pistaferri, and Alessandra Voena. 2018. "Marriage, Labor Supply and the Social Safety Net." *Unpublished manuscript* .
- Lundberg, Shelly J, Robert A Pollak, and Terence J Wales. 1997. "Do husbands and wives pool their resources? Evidence from the United Kingdom child benefit." *Journal of Human Resources* :463–480.
- Mazzocco, Maurizio, Claudia Ruiz, and Shintaro Yamaguchi. 2013. "Labor supply, wealth dynamics, and marriage decisions." *Unpublished manuscript* .
- Pendakur, Krishna. 2018. "Welfare analysis when people are different." *Canadian Journal of Economics/Revue canadienne d'économique* 51 (2):321–360.
- Piketty, Thomas and Emmanuel Saez. 2013. "Optimal labor income taxation." In *Handbook of Public Economics*, vol. 5. Elsevier, 391–474.
- Reynoso, Ana. 2019. "The impact of divorce laws on the equilibrium in the marriage market." *Unpublished manuscript* .
- Santaeulàlia-Llopis, Raül and Yu Zheng. 2017. "Why Is Food Consumption Inequality Underestimated? A Story of Vices and Children." *Unpublished manuscript* .
- Shephard, Andrew. 2019. "Marriage market dynamics, gender, and the age gap." *Unpublished manuscript* .

Appendix

Appendix A: Computational details

Solving the model requires finding a fixed point between the *expected* distributions of singles, which individuals use to compute the expected values over the future, and the *actual* distributions of singles, which depend on the marriage and divorce choices that individuals make. The equilibrium essentially requires rational expectations, so that the beliefs about the future coincide with the actual distributions of singles. The basic algorithm is to

- start with a guess of the expected distributions of the characteristics of singles at each age $\Lambda_{t,g}(\epsilon, H, \alpha, A)$
- solve the life-cycle problem for the value functions given these expected distributions
- simulate the model, where individuals make choices about consumption and time use, marriage and divorce given the expected value functions from the previous step
- compute the implied distributions based on the simulations
- update the expected distributions and repeat until convergence

Note that little can be said about the theoretical equilibrium properties, regarding existence and uniqueness, as is standard in search models of marriage (also see [Guvenen and Rendall \(2015\)](#) or [Greenwood et al. \(2016\)](#)). In numerical checks, this does not pose an issue given the calibrated parameter values.

The life-cycle problem

Given a guess for the expected distribution of singles at each age, the life-cycle problem can be solved via backwards-induction starting in the terminal period T . The resulting value function has the interpretation as the *expected* value function given expectations about the future availability of singles. Both the assets and the Pareto weights are treated as continuous state variables. Values outside of the corresponding grids are interpolated linearly. The distributions of singles at each age are discretized on a fine grid in the asset dimension. Note that temporary wage shocks do not introduce separate state variables, which facilitates including them without greatly increasing the size of the state space. Since the temporary shocks are i.i.d. and do not influence continuation values, for example a positive wage shock is equivalent to an increase in assets (conditional on labor supply), so that the temporary shocks can be evaluated on the same grid that is also used to evaluate different asset levels.

Simulation and updating the distributions

Having computed the value function, one needs to update the guess of the distributions of singles, which is done via simulation. In the simulation, potential partners are always drawn from the

actual distribution of singles, which can be computed from the currently available singles in the simulation (given the assumption of meetings only occurring within cohorts). The guess for the *expected* distributions of singles enters (only) through the value functions, which individuals use for example to compare the value from getting married to the value of staying single. In the end, the guess for the distributions is updated (with a weight κ on the old distributions).

Implementation

The model is implemented with Python and Numba and solved on a computing cluster. The parallelization is based on MPI. To get good performance for this type of model, it is helpful to use MPI only for inter-node communication and use shared-memory parallelism (such as Numba's *prange*) within a node. This limits communication time and allows to have a full node as a master MPI rank with sufficient memory for the simulations.

Appendix B: Results for alternative inequality measures

For robustness, I also conducted the decompositions for two alternative inequality measures that are commonly used. For the purpose of this paper, the Theil index and the Mean Logarithmic Deviation (*TI* and *MLD* in the following) are most suitable, because they can be decomposed into within and between components. As a result, the same decomposition formulas that were used in the main text can be applied. The two indices are defined as:

$$TI = \frac{1}{N} \sum_{i=1}^N \frac{X_i}{\bar{X}} \log \left(\frac{X_i}{\bar{X}} \right)$$

$$MLD = \frac{1}{N} \sum_{i=1}^N \log \left(\frac{\bar{X}}{X_i} \right)$$

The tables replicates the main results for these two inequality measures. Since they require the variable to be positive, I report the results for the equivalence scale (instead of the utility measure) and private consumption. The results are in tables [A1](#) and [A2](#). The first row in each section shows the decomposition of the level of the inequality measure and the second row shows the decomposition of the change. The main conclusions from the tables are similar as when using the variance. In addition, the differences between the Theil index and the MLD are small.

TABLE A1: Decomposition of Mean Logarithmic Deviation

	Total	Singles (%)	Within couples (%)	Between couples (%)	Betw. sin. mar. (%)
(a) c_1					
MLD ($\psi_2 = 0.15$)	4.7	21.7	23.6	54.3	0.4
Change ($\psi_2 = 0.21$)	-11.2	0.15	0.23	0.62	-0
(b) Log eq. scale					
MLD ($\psi_2 = 0.15$)	9.8	13.5	6.5	69.7	10.2
Change ($\psi_2 = 0.21$)	-4.9	0.27	0.1	0.61	0.02
(c) Log eq. scale (LT)					
MLD ($\psi_2 = 0.15$)	6.9	16.9	8.5	74.5	0.1
Change ($\psi_2 = 0.21$)	-8	0.22	0.09	0.69	-0

Notes: This table replicates the decomposition for the MLD.

TABLE A2: Decomposition of Theil index

	Total	Singles (%)	Within couples (%)	Between couples (%)	Betw. sin. mar. (%)
(a) c_i					
Theil index ($\psi_2 = 0.15$)	4.7	21.1	23.3	55.1	0.4
Change ($\psi_2 = 0.21$)	-11.4	0.16	0.24	0.61	-0
(b) Log eq. scale					
Theil index ($\psi_2 = 0.15$)	9.8	9.7	7.5	73.3	9.5
Change ($\psi_2 = 0.21$)	-5.4	0.19	0.13	0.67	0.02
(c) Log eq. scale (LT)					
Theil index ($\psi_2 = 0.15$)	7.2	16.8	8.3	74.8	0.1
Change ($\psi_2 = 0.21$)	-8.4	0.23	0.1	0.68	-0

Notes: This table replicates the decomposition for the Theil index.

Appendix C: Details of the decomposition formula

The expression for the variance was given by:

$$V(X) = p_s V^S + p_m V^{M,W} + p_m V^{M,B} + V^{BMS}$$

The variance between couples and singles can be split into two parts:

$$\begin{aligned} V^{BMS} &= p_s V^{BMS,S} + p_m V^{BMS,M} \\ V^{BMS,S} &= (E(X| \text{Single}) - E(X))^2 \\ V^{BMS,M} &= (E(X| \text{Married}) - E(X))^2 \end{aligned}$$

Overall, this leads to the following decomposition of the growth rate:

$$\hat{V}(X) = \omega_0 \hat{p}_m + \omega_1 \hat{V}^S + \omega_2 \hat{V}^{M,B} + \omega_3 \hat{V}^{M,W} + \omega_4 \hat{V}^{BMS,S} + \omega_5 \hat{V}^{BMS,M} + \hat{R}$$

The weights ω_i and \hat{R} and residual are the following:

$$\begin{aligned} \omega_0 &= \frac{-p_m V^s + p_m V^{M,B} + p_m V^{M,W} - p_m V^{BMS,S} + p_m V^{BMS,M}}{V} \\ \omega_1 &= \frac{p_s V^S}{V} \\ \omega_2 &= \frac{p_m V^{M,B}}{V} \\ \omega_3 &= \frac{p_m V^{M,W}}{V} \\ \omega_4 &= \frac{p_s V^{MBS,1}}{V} \\ \omega_5 &= \frac{p_m V^{MBS,2}}{V} \\ \hat{R} &= \hat{p}_m \hat{V}^S \cdot \frac{-p_m V^s}{V} + \hat{p}_m \hat{V}^{M,B} \cdot \frac{p_m V^{M,B}}{V} + \hat{p}_m \hat{V}^{M,W} \cdot \frac{p_m V^{M,W}}{V} \\ &+ \hat{p}_m \hat{V}^{BMS,S} \cdot \frac{p_m V^{BMS,S}}{V} + \hat{p}_m \hat{V}^{BMS,M} \cdot \frac{p_m V^{BMS,M}}{V} \end{aligned}$$

Appendix D: Calibration of the wage process

Since the stochastic process for wages is exogenous, it can be estimated separately from the rest of the model. I estimate the parameters of the process by matching moments of the process to data moments from the Dutch Socio-Economic Panel. Since the process is assumed to be the same for men and women, a convenient way to address potential selection problems is to estimate the process on data for men only. The parameters that need to be calibrated are the persistent and variance of the wage process (which is discretized according to the Rouwenhorst method), the variance of the temporary earnings shock, and the mean and variance of the ability distribution. The ability distribution is a discretized log-normal distribution. To separately pin down variability over time (through shocks) and persistent heterogeneity, I include moments on the variance of wage changes as well as the cross-sectional variance of wages. In addition, I include two autocorrelation moments to pin down the persistence parameter of the process and the variance of the temporary shock. The fit of the wage process and the parameters are summarized in tables A3 and A4.

TABLE A3: Wage moments - fit

	Model	Data
$Var(\log(w_{t+1,i}) - \log(w_{t,i}))$	0.03	0.02
$Corr(\log(w_{t+1,i}), \log(w_{t,i}))$	0.93	0.86
$Corr(\log(w_{t+2,i}), \log(w_{t,i}))$	0.89	0.75
$E(\log(w_{t,i}))$	2.78	3.03
$Var(\log(w_{t,i}))$	0.21	0.15
$q_{\log(w_{t,i})}^{10}$	2.19	2.60
$q_{\log(w_{t,i})}^{90}$	3.43	3.49

Notes: This table compares the moments of the wage process to the data.

TABLE A4: Parameters of the wage process

	Value
σ_w	0.41
ρ_w	0.94
Variance of temp. shock	0.05
Mean of persistent heterogeneity	1.10
Variance of persistent heterogeneity	0.09

Notes: This table shows the calibrated wage parameters.

The process for women is identical to the process for men except for the parameter for the gender wage gap. This parameter is set to 0.83, which reproduces the difference in mean earnings between men and women in the data. Note that this abstracts from potential selection issues as well as from endogenous human capital accumulation, which could be included in a richer specification.⁴⁶

⁴⁶With human capital accumulation, the gender wage gap would be endogenous.

Appendix E: Additional tables and figures

TABLE A5: Composition of inequality under the two policy regimes

	$\psi_2 = 0.15$	$\psi_2 = 0.21$
(b) $\log(c_i)$		
Within couples (% of total)	23.9	23.9
Between couples (% of total)	54.3	53.3
Singles (% of total)	21.4	22.2
Betw. sin. and mar (% of total)	0.4	0.6
(a) $u(c_{it}, C_{it}, l_{it}, D_{it})$		
Within couples (% of total)	6.2	6.1
Between couples (% of total)	67.9	68.6
Singles (% of total)	17.3	16.4
Betw. sin. and mar (% of total)	8.6	8.8

TABLE A6: Variance decomposition - Reduction of progressivity to $\psi_2 = 0$

	$\log(\mathbf{c}_i)$	$\mathbf{u}(\mathbf{c}_i, \mathbf{C}_i, \mathbf{l}_i, \mathbf{D}_i)$	\mathbf{u}_{it}^{LT}
Change (%)	27.8	12	17.5
Decomposition of change			
Within couples (%)	25.54	8.33	9.71
Between couples (%)	61.87	53.33	70.86
Singles (%)	13.67	26.67	21.14
Betw. sin. and mar. (%)	-1.08	10.83	-0.57
Single probability (%)	0	0	-0.57

Notes: This table shows the effect of the policy change on the variance of private consumption, the per-period utility and life-time utility. The first row shows the total change. The other rows decompose this total change in inequality and add up to 100%. Compared to the results from table 5, the table illustrates that increasing or reducing progressivity has relatively symmetric effects.