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

Tim Obermeier
University of Mannheim

Job Market Paper

[Link to most recent version]

November 12, 2018

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. 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 reform decreases intra-household inequality in private consumption by 10.7%, while it decreases inequality across households by 15.4%. The reduction of intra-household inequality accounts for 16.5% of the reduction in inequality in private consumption due to the reform, and 6.9% of the reduction in inequality in welfare from private and public consumption, leisure and home production. Marriage rates increase especially for men with relatively low ability.

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, 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. Financial support through CRC TR 224 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* households and paid little attention to the question how inequality between spouses *within* households 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 (2017)). 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, this paper develops an equilibrium search and matching model with intra-household bargaining, labor supply and savings and calibrates the model to data from the Netherlands. 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. Personal consumption and leisure are private goods. In the model, 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 could afford a higher living standard as 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, and thereby raises her share of consumption and leisure in marriage.

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. When meeting on the marriage market, individuals take both the economic benefits and the non-economic quality of the relationship into account. They weigh the benefits from getting married against the value of waiting to meet someone else. 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 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.

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

I calibrate the model to Dutch survey data on intra-household allocations and labor market outcomes. As the majority of OECD countries, the Netherlands has a system of individual taxation, in which spouses are taxed based on their individual earnings.[2] The calibrated model is used to study the effects of budget-balanced increases in progressivity, in which the average tax rate falls for low incomes and rises for higher incomes. Before turning to the policy analysis, I first decompose 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 20.1% of the total variance of private consumption, singles for 30.3% and the between-couple component for 49.5%. 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 5.4% 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. The experiment increases progressivity by 40% and decreases the variance of private consumption by 12.7%.[3] The focus of the analysis is to decompose the total change in inequality due to the reform in terms of inequality within and between couples and of singles. The reform reduces both the within- and the between- household variance. The contribution of the within-household component to the total change in inequality is 16.5% for private consumption and 6.9% for the utility measure. Thus, in terms of utility from private and public goods, 6.9% of the total change in inequality would be ignored by abstracting from intra-household inequality. The increase in progressivity affects inequality both bv changing bargaining positions and by changing marriage rates. Since there are economies of scale in consumption and home production, married individuals have a higher living standard than singles. Changes in marriage rates, due to the reform, therefore contribute to the effect on inequality. In the model, singles are more likely to meet others from the same quartile of the ability distribution. Especially men who have low earnings relative to the average of their quartile marry at a lower rate, compared to other men from the same ability quartile. Women with relatively low wages marry at a *higher* rate. This is due to gender asymmetries with respect to wages and home production. Women are more likely to reduce

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

[3]The change in progressivity refers to the progressivity parameter of the tax function used in Heathcote, Storesletten, and Violante (2017) or Holter, Krueger, and Stepanchuk (2014).

their market hours when having kids and contribute less to the total income of the household. The increase in tax progressivity increases marriage rates and men with the relatively low wage in their quartile are more likely to get married.

To study to what extent the marriage market channels (marriage and bargaining) contribute to the total welfare effect of the reform, I further consider changes in life-time utility. I isolate the role of the marriage market by computing the welfare effect of the reform for a 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 this case to the full model, where bargaining weights and marriage decisions adjust. The effect of the marriage market, defined as the difference between the full model and the model without marriage market adjustments, differs along the ability distribution and mostly ranges between 64% and 3% of the total effect of the reform.

The main contribution of this paper is that it analyzes how tax progressivity affects inequality through its impact on intra-household inequality and the marriage market. The literature on tax progressivity typically focuses on 'one-earner' models, in which the presence of a spouse is not explicitly modeled, and thus does not allow for intra-household inequality. In the macroeconomic literature, Conesa and Krueger (2006) and Heathcote, Storesletten, and Violante (2017) have studied 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)), but typically take intra-household allocations and marriage outcomes. Notable exceptions are Alesina, Ichino, and Karabarbounis (2011) and Bastani (2013), who study how gender-based taxation, rather than tax progressivity, affects intra-household bargaining. Gayle and Shephard (2018) 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. My paper contributes by focusing on progressivity and on decomposing its effect on inequality in terms of within- and between household inequality. In addition, I analyze a dynamic marriage market, which leads to different mechanisms.

My research further relates to papers have used dynamic bargaining models to study the effects of welfare time limits (Low et al. (2017)) or tax credits (Mazzocco, Ruiz, and Yamaguchi (2013)). My paper analyzes a different policy and further differs by modeling the marriage market equilibrium. Equilibrium models of marriage and intra-household allocations are still relatively rare (see e.g. Goussé, Jacquemet, and Robin (2017), Chiappori, Dias, and Meghir (2018) or Reynoso (2017) 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 my paper on the effects of changing bargaining positions 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 (2017) analyze Japanese panel data and relate allocations to differences in relative wages. Santaeulàlia-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 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 and 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 proceeds 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 e.g. trips, housing expenditure, 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. Panel (a) of figure 1 shows the ratio of the wife's private (assignable) consumption relative to total private consumption.

The figure illustrates that spouses often consume different amounts of private goods. In a within-between decomposition of the variance of private consumption, the the within-couple component has 51.1% the size of the between-couple component.[4] Panel (b) shows a similar graph

[4]The within-between decompositions are described in more detail in section 5.1.

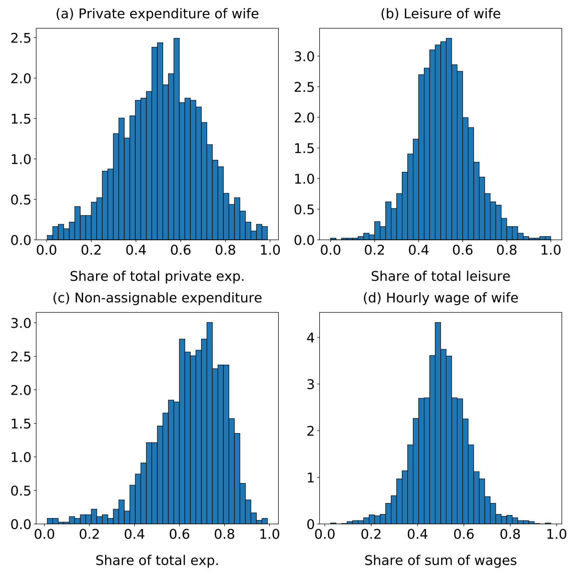

FIGURE 1

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.

for leisure, which also features substantial dispersion. Panel (c) depicts the ratio of non-assignable expenditure to total expenditure and shows that accounting for public consumption is important, since the non-assignable expenditure is an important category. Finally, panel (d) shows the ratio of the hourly wage of the wife relative to the sum of wages, for all couples in which both partners work. There are many couples in which wage rates are relatively unequal, even though education levels of spouses are typically quite similar. This is e.g. due to the fact that there is still substantial variation across wage levels across occupations within an education level.[5]

Lise and Yamada (2017) 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.

[5]The focus is on wage rates, rather than earnings, since the wage rate can be a better proxy for would a spouse *could* earn if they were single, which corresponds more closely to the outside option in a collective model.

3 Model

3.1 Overview

The model is an equilibrium search model of marriage, divorce, labor supply and savings. There are overlapping generations of men and women that live for $T = 21$ periods, each period representing 3 years. Individuals participate in the labor market until they reach the retirement age T_R. They start their life single and can meet other singles from their cohort on the marriage market in each period. Individuals are heterogenous in their labor market ability and face permanent and transitory shocks to their earnings potential over their lifetime, which they can partially insure against using savings.

The most important model ingredients concern how wage inequality translates into inequality in consumption and welfare, in terms of utility from leisure and home production. Within marriage, not all goods are shared equally between spouses. There is a private consumption good and leisure is private good. As a result, one spouse may have a higher level of consumption and leisure than the other. Intra-household allocations are modeled as a limited commitment contract (Mazzocco, Ruiz, and Yamaguchi (2013), Voena (2015)). The central variable that determines how much utility a married indiviual is allocated is the ’decision weight’ of a spouse. This weight is initially determined at the time of marriage. Spouses can commit to this allocation until the participation constraint of one spouse binds (i.e. they would prefer to divorce). In this case, the allocation can be renegotiated. The initial bargaining weight is determined as the solution to a Nash bargaining problem. The outside option in this bargaining problem is to stay single for another period and be matched with another individual. After marriage, the outside option is to divorce and re-enter the marriage market.

Marriage leads to economies of scale, since there is a joint consumption good and spouses can pool their home production time. Couples where both partners have a high income can afford a high level of public consumption. A final economic benefit of marriage is risk-sharing. In addition, there is a non-economic match quality component. Individuals make decisions about marriage and divorce taking both the economic and non-economic benefits into account. Due to the economic benefits from marriage, marriage and divorce have important economic implications. Singles meet other singles from their cohort on the marriage market. The (steady-state) marriage market equilibrium requires that individuals meet spouses from the available pool of singles at each age, and that their expectations about future partners are consistent with the actual pool of singles.

In the following, the model will be described in more detail. Section 3.7 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 consumption 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:

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

The parameter z controls the substitutability between the two time inputs. The home production technology introduces a gender asymmetry, since the productivity of the time inputs differs. The parameter η_m corresponds to $1 - \eta_f$. Individuals discount the future at rate β and have access to a risk-free asset that yields an interest rate R.

3.3 The labor market and the government

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

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

κ_i represents the ability of the indivdual. κ_g^G refers to the gender wage gap. There are two different wage shocks.[6] $\epsilon_{i,t}$ 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

[6] Note 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 type distribution with 12 types in order 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.

variance σ_T.[7] 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 modeled via a tax function, using the functional form from Heathcote, Storesletten, and Violante (2017):

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

The parameter ψ_2 refers to the progressivity of the system and ψ_1 is a level parameter. Married couples are taxed individually in the Netherlands, so that the tax liability is calculated for each spouse separately.[8] For simplicity, this specification abstracts from aspects like transferable deductions between spouses, which would introduce a small degree of dependence between partners.[9] 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 exogenous spending requirement $\bar{G}$:

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

Here, B indicates whether an individual gets the social assistance benefit. The government redistributes in the sense that it can choose whether to raise revenue mainly from higher-income individuals or more evenly across the distribution.

3.4 Singles

Singles decide on consumption, savings and the time spent for home production, leisure and labor supply. Their value function and budget constraint are:

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

The current asset level is denoted as $A_{i,t}$, while A' represents savings for the next period. $\text{EV}^S_{i,t}(A')$ is the continuation value and also takes the outcomes of the marital matching process into account, as will be described below. State variables, or objects that depend on state variables, are indexed

[7] The 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.

[8] The 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. Note that this specification models the tax and not the transfer system, as $T(y)$ is restricted to be positive. This is due to the fact that in reality transfers are typically means-tested on the family level.

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

by the individual and period. When two individuals meet or are married, the individuals are indexed by their gender $g \in \{f, m\}$, rather than their individual index.

In the first period of life, singles draw an initial 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 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.[10]

In the beginning of each period, individuals observe the realization of the persistent productivity shock and whether the children, if there are any, grow up. The probability of participating in the marriage market is given by $M(t)$, which declines with age. Individuals who enter the marriage market meet a potential partner from the pool of available singles. The characteristics of the potential spouse are their ability, productivity, assets and children. The probability of meeting someone with these characteristics, $\Lambda_{t,g}(\epsilon, \alpha, A, b)$, is an equilibrium object, because it depends on which individuals got married or divorced in previous periods 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 θ is drawn from a $N(0, \sigma_\theta)$ distribution for each match, which captures the non-economic quality of the potential marriage.

At the time of the meeting, both individuals observe the match quality draw and the persistent characteristics of the potential partner. To simplify the computations, I assume that the temporary earnings shocks are realized *after* the marriage stage.[11] The individuals have to decide whether to get married and, if so, how to allocate decision power in marriage. Decision power is summarized by a Pareto weight $\lambda \in [0, 1]$, which is discussed in more detail in the next section. The decision weight determines how much utility an individual receives in marriage, which is denoted as $\mathbb{E}_w V^M_{g,t}(\lambda, \theta_{i,t}, \omega_{f,t}, \omega_{m,t})$. The expectation $\mathbb{E}_w$ refers to the expectation over the temporary wage shocks. The value of rejecting the match, on the other hand, is the expected value of staying single for one period and drawing another potential partner in the next period. This value is the 'outside option' for bargaining over the Pareto weight.

The potential couple has to determine if there is a range of weights $\lambda \in (0, 1)$ so that both individuals prefer to get married over staying single ($\mathbb{E}_w V^M_{g,t}(\lambda, \omega_{f,t}, \omega_{m,t}) \geq \mathbb{E}_w V^S_{g,t}(\omega_{g,t})$), given individual characteristics $\omega_{g,t}$. If this is the case, the Pareto weight of the couple results from Nash bargaining:

$$\tilde{\lambda} = \text{argmax}_\lambda (\mathbb{E}_w V^M_{f,t}(\lambda, \omega_{f,t}, \omega_{m,t}) - \mathbb{E}_w V^S_{f,t}(\omega_{f,t}))(\mathbb{E}_w V^M_{m,t}(\lambda, \omega_{f,t}, \omega_{m,t}) - \mathbb{E}_w V^S_{m,t}(\omega_{m,t}))$$

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

[11]Since these shocks are relatively small and temporary, they would have a small impact on marriage decisions and bargaining. The assumption simplifies computing the integral over all matches that can occur in a period, since one needs to compute the bargaining solution for each possible combination of state variables that can occur. In addition, the assumption smoothes out the Pareto frontier (at the time of bargaining), which would otherwise have discontinuities due to discrete labor supply.

If there is no Pareto weight such that both individuals prefer to get married, the match is rejected and they stay single. Taken together, the continuation value for singles is the expected utility over all matches that may occur and the corresponding marriage decision in each case:

$$\mathrm{EV}^S_{g,t}(A') = \int_{\omega^M=(\omega_{f,t+1},\omega_{m,t+1},\theta)} M(\omega^M_{t+1}) \cdot \mathbb{E}_w V^m_{g,t+1}(\lambda(\omega^M_{t+1})) + (1 - M(\omega^M_{t+1})) \cdot \mathbb{E}_w V^S_{g,t+1} \,\mathrm{d}F(\omega^M_{t+1})$$

ω^M_t is a vector containing the characteristics of both individuals and the match quality realization. $M(\omega^M_t)$ is an indicator variable that takes the value of 1 if marriage occurs given these match characteristics and $\lambda(\omega^M_t)$ is the Pareto weight that is chosen in this case.

3.5 Couples

Married couples decide on consumption, labor supply and home production time. The utility of a married individual, conditional on these choices is:

$$U_{i,t}(c, C, l, D, A') = u_{i,t}(c, C, l, D) + \theta_{i,t} + \beta \mathrm{EU}_{i,t}(\lambda_{i,t}, A')$$

θ_t denotes the current match quality of the couples. Match quality changes over time according to an $AR(1)$ process with persistence ρ_θ and variance σ_θ. $\mathrm{EU}_{i,t}(\lambda_{i,t}, A')$ 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 $p^{b,g}$.[12] 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 level of match quality (θ) and the presence of children (b):

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

In the beginning of the marriage, the Pareto weight of the couple is given by the value that was determined on the marriage market. The weight $\lambda_{i,t}$ refers to the weight of the woman, the weight of the man is given by $1 - \lambda_{i,t}$. The household maximizes the weighted sum of utilities of the spouses:

$$\max_{\mathscr{C}=(c_f,c_m,C,A',h_f,h_m,d_f,d_m,l_f,l_m)} \lambda_{i,t} U^f_{i,t}(\mathscr{C}) + (1 - \lambda_{i,t}) U^m_{i,t}(\mathscr{C})$$

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

The constraints are given by the budget and time constraints:

$$
\begin{aligned}
c_f + c_m + C + A' &= w_{f,t}h_f + w_{m,t}h_m - T(w_{f,t}h_f) - T(w_{m,t}h_m) + RA_{i,t} \\
h_f + d_f + l_f &= 1 \\
h_m + d_m + l_m &= 1
\end{aligned}
$$

The utility levels evaluated for the optimal choices are denoted as $U^*_{f,t}(\lambda_{i,t}, \omega^M_{i,t})$ and $U^*_{m,t}(\lambda_{i,t}, \omega^M_{i,t})$, given λ and the state variables of the couple ω_c.

Spouses can unilaterally file for divorce in the beginning of a period and decision-making is subject to limited commitment (Mazzocco, Ruiz, and Yamaguchi (2013), Voena (2015)). This means that the Pareto weight of the couple remains the same over time until one spouse would prefer divorce over staying married at this weight. The participation constraint of each spouse is given by:

$$
U^*_{g,t}(\lambda_{i,t}, \omega_c) \geq V^S_{g,t}(A_{i,t}/2, \alpha_{g,t}, \epsilon_{g,t}, b_{i,t})
$$

Assets are split equally in divorce and children stay with the mother, who also takes them into the next marriage.[13] If the participation constraint binds for both spouses, divorce occurs. If only one participation constraint binds, the Pareto weight of this spouse is (minimally) adjusted to make him or her indifferent between staying and divorcing. If this is feasible and the other spouse, whose participation constraint was not violated, also prefers to stay married with the new decision weight, the weight adjusts. If not, the couple divorces.

The interpretation of limited commitment is that risk-sharing in marriage is imperfect. Match quality and wages fluctuate over time. In particular, changes in wages due to permanent productivity shocks can lead to a change in bargaining power. 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.[14]

3.6 Marriage market equilibrium

3.6.1 Meeting technology

In the beginning of each period, individuals can participate in the marriage market and may meet other singles from the distribution of available singles. Individuals only meet potential partners from the same cohort.[15] The probability of participating in the marriage market, $M(t)$, de-

[13] The 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.

[14] Evidence from Lise and Yamada (2017) supports these models of decision-making, as they find that decision power changes infrequently and more often before divorce and for big shocks.

[15] This assumption is a good approximation of the data. Allowing for meetings between cohorts would significantly enlarge the state space, since it would require keeping track of the age of both spouses, instead of only the common

clines with age, and models in a simple way the fact that it may be harder to meet someone in later life, when there are fewer singles. Given a meeting, individuals have to decide whether to get married or to keep searching.

Meetings are assortative based on the ability type, which captures that individuals meet similar potential spouses at a higher rate (e.g. at the workplace or in their social circle).[16] This allows the model to better fit patterns of assortative mating in the data. For example, with perfectly random meetings there would be few meetings between two individuals of the highest type, since their share in the population is low, and some high types would not even meet another high time within their lifetime. I make the following assumptions, which introduce assortative meeting, while ensuring that meetings are consistent with the available pool of singles. With 12 types in total, three adjacent types form an 'ability group' (e.g. types 1, 2 and 3) and are more likely to meet each other, as opposed to individuals from other ability groups. This has the intuitive appeal that individuals are more likely to meet others from their social environment (the ability group), while some individuals are below or above the average within the ability group.

More concretely, in the beginning of each period, m_{am} percent of singles from each gender and ability group are taken from the current pool of singles and randomly matched with individuals within the same ability group, as long as this is possible. Conditional on meeting a potential partner with a given type, the remaining characteristics (the permanent wage shock, assets and children) are drawn from the distribution of characteristics of that type. It can be the case that some individuals are unmatched at this stage. This occurs when the size of the ability cells differs by gender, since then taking m_{am} percent of each cell leads to different masses. These individuals return to the pool of singles. In a second stage, the remaining individuals are matched randomly with each other. The parameter m_{am} can be matched to the data to reproduce the empirical degree of wage differences within couples.

These assumptions ensure that meetings are consistent with the available pool of singles of each gender. Given measure f_i of type i females and m_j of type j males, there is a measure $\mu_{i,j}$ of meetings between the two ability types. With perfectly random meetings, the number of meetings would simply be $\mu_{ij} = f_i \frac{m_j}{\sum m_j}$. Consistency requires the following three conditions:

$$\begin{aligned} \sum_i \sum_j \mu_{i,j} &= \sum_i f_i \\ \sum_j \mu_{i,j} &= f_i \\ \sum_i \mu_{i,j} &= m_j \end{aligned}$$

The first condition implies that the total mass of meeting equals the mass of individuals. Note that the total mass of single males and females is always equal by assumption (there is no imbalance in the sex ratio, so that $\sum_i f_i = \sum_j m_j$). The other two conditions require that the total mass of

age.

[16]See Guvenen and Rendall (2015) for related assumptions on assortative meeting.

females that men meet must be equal to the actual mass of females (and vice versa).

3.6.2 Definition of equilibrium

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

1. the policy functions $(A, c, C, h, l, d) = P^S_{g,t}(\omega^S)$ solve the problem of singles
2. the policy functions $(A, c_f, c_m, C, h_f, h_m, l_f, l_m, d_f, d_m) = P^M_t(\omega^M)$ 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)

3.6.3 Discussion and computation

The marriage market equilibrium requires that meetings are consistent with the available pool of singles and that individuals have correct expectations over the future pool of singles. These expectations are important when individuals decide whether to get married or divorced, since the value of being single depends on the distribution of spouses one will meet in future periods. In addition, the expectations over the distributions of spouses can also influence labor supply and savings decisions. Equilibria are 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.

3.7 How does tax progressivity affect intra-household inequality?

To illustrate the mechanisms of the model, this section discusses an example to show how bargaining weights are determined (also see the discussion in Knowles (2012)) 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 individuals. 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 bargaining weight. The dot shows the bargaining outcome given the outside options.

The graph on the right side shows the case where the government redistributes such that 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. This leads to the following first-order condition:

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

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. The first-order condition shows the main determinants of the Pareto weight. 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 bargaining 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 bargaining 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. In the calibrated model, the slopes of the Pareto frontiers are such that reducing utility of the more well-off individual by 1, relative to the bargaining outcome, increases the utility of their partner by 2.04 for the median couple (the 25 and 75% percentiles are 1.51 and 2.99).

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 bargaining weights.[17] 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 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.

[17]To see to what extent differences in future marriage market outcomes determine bargaining weights, I compared the weights from the calibrated model to the weights that would be obtained if the outside option was to stay single forever. This is shown in figure A1 in appendix D. The figure indicates that the values of being single alone over- or underpredicts the Pareto weight by up to 0.2 in most cases.

4 Calibration

4.1 Calibrated parameters

A number of parameters are set externally. The coefficient of risk aversion is set to 2, which is a standard value in the literature. The discount factor is 0.913, which implies a yearly discount factor of 0.97 (one period are three years). The interest rate is set such that $R = 1/\beta$ holds. The preference parameter for the private good, α_c, is set to 0.176, which implies that households spend roughly 61% of their consumption expenditure on private goods.[18] 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. Since wages are exogenous in the model, the wage process can also be set externally. This is done using data from the Dutch Socioeconomic Panel, due to its larger sample size. The wage process is calibrated to match changes in wages within individuals. 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 scale parameter for assortative matching.[19] Regarding the moments, I match the share of individuals that is currently married or cohabiting at age 20 and 35, the share that ever married or cohabiting by age 45, the share that ever experienced a divorce or separation by ages 35 and 45.[20] These moments are constructed based on the Dutch Kinship Survey, which provides restrospective 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. 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 weighted distance between the model and data moments. To avoid computing the equilibrium distributions for each set of trial parameters - in particular for those cases, where the trial parameters are far from the data -, I start with a reasonable initial guess for the distributions, solve and simulate the model only once for each parameter vector, then update the distributions with a very low weight on the new distributions, and proceed

[18] In practice, the exact share varies a little, since it also depends 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.

[19] Recall 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.

[20] In 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 cohabitating couples that are not captured by the model (relating e.g. to the division of assets upon divorce).

TABLE 1: Calibrated parameters

Parameter	Value
Scale of leisure preference	0.48
Scale of home good (no kids)	0.12
Scale of home good (kids)	0.63
Variance of match quality	0.98
Persistence of match quality	0.84
Meeting rate in last period	0.05
Scale for assortative meeting	0.64
Home productivity (women)	0.57

Notes: This table shows the calibrated parameters.

with the next parameter vector. 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. 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.

4.2 Model implications and comparison to data

Before turning to the policy analysis, this section first shows how the model compares to the data, in terms of inequality within and between households. Figure 3 first shows the distribution of household income in the economy, including both (working age) singles and married couples and shows the income distribution from the data. The figure is constructed using data from the Income Panel Study (*Inkomenspanelonderzoek*), which is based on tax records on contains detailed information on income for a large sample of individuals. [21] The horizontal axis shows yearly labor income. The graph illustrates that the model features substantial inequality across households and compares well to the data.

Figure 3 also shows the relative private consumption and leisure within couples. The graphs show 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. They compare well to the data, which exhibits somewhat more dispersion than than the model. To illustrate the Pareto weights underlying these allocations, table 3 shows the mean Pareto weight for each combination of ability groups. Group 1 is the highest.

[21] The publicly available data is a repeated cross-section.

TABLE 2: Model fit

Moment	Model	Data
Currently married, age 20	0.26	0.24
Currently married, age 35	0.74	0.79
Ever married, age 45	0.89	0.92
Ever divorced, age 35	0.25	0.30
Ever divorced, age 45	0.36	0.33
Work hours, women, couples without kids	0.38	0.35
Work hours, women, couples with kids	0.22	0.20
Home hours, women, couples without kids	0.17	0.21
Home hours, women, couples with kids	0.40	0.48
Relative home hours	0.64	0.61
Share of within-couple wage variance	0.48	0.40

Notes: This table summarizes the fit of the model.

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 (4) and the husband in the highest (1), the mean Pareto weight is 0.13.

For policy analysis, a relevant implication of the model is is how strongly the Pareto weight varies with changes in relative wages. This can be compared to the empirical estimate from Lise and Yamada (2017), who 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 1.7%, indicating that the implied elasticity of the Pareto weight seems reasonable and compares well to empirical studies.

TABLE 3: Pareto weights

	M - Group 1	M - Group 2	M - Group 3	M - Group 4
F - Group 1	0.49	0.58	0.70	0.87
F - Group 2	0.40	0.48	0.62	0.82
F - Group 3	0.28	0.35	0.48	0.71
F - Group 4	0.13	0.17	0.27	0.47

Notes: The table shows the Pareto weights for couples across ability groups. Group 1 is the highest, group 4 for the lowest.

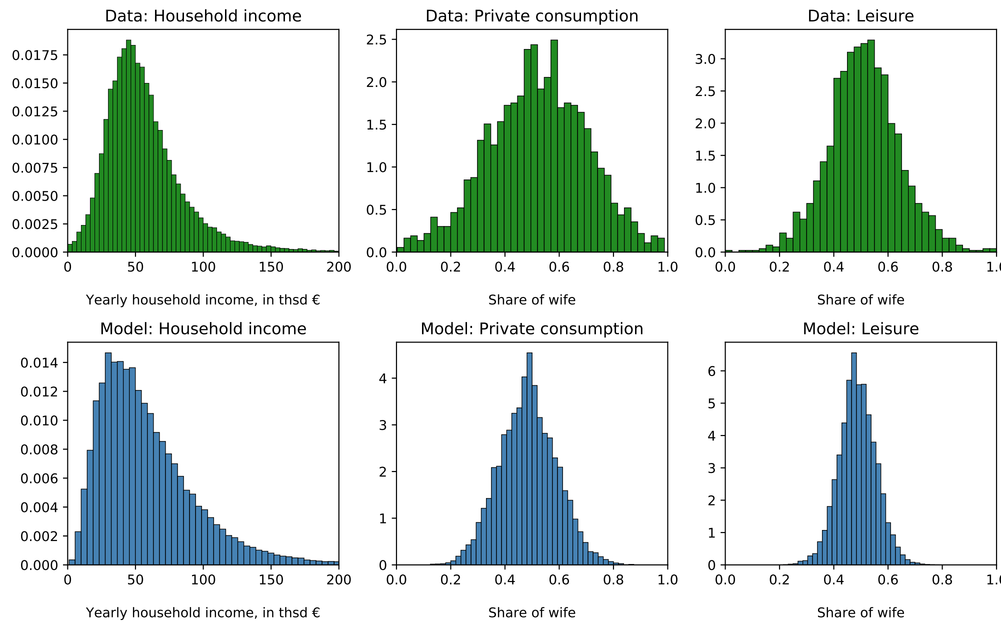

FIGURE 3

Notes: The figure shows the distribution of household income and relative 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 (e.g. $\frac{c_f}{c_f+c_m}$*)*

5 Results

In this section, I first discuss cross-sectional inequality in the calibrated model, given the current tax schedule, in more detail. I then describe the policy experiment - a hypothetical increase in progressivity - and show how the reform affects inequality within- and between households and marriage market outcomes. Finally, I turn to alternative welfare measures, robustness and life-time welfare.

5.1 Decomposition of inequality

The focus of the inequality analysis is on consumption and welfare.[22] I first consider inequality in private consumption, since this allows to study the role of intra-household bargaining most directly. The model further allows to study inequality in (per-period) *utility* from private and public consumption, leisure and home production. The main advantage of this measure is that it captures all private and public goods that are available to individuals and is therefore a more complete indicator of welfare. In particular, the allocation of leisure is an important part of intra-household bargaining and focusing on consumption only would ignore this aspect. The utility

[22]This relates to a large literature on consumption inequality (e.g. Lise and Seitz (2011), Krueger and Perri (2006), Blundell and Preston (1998)). In addition, there a few recent papers that include time-use into welfare measures (see e.g. or Cherchye et al. (2018) or Chiappori and Meghir (2015) for a theoretical discussion).

measure further takes public goods in terms of public consumption and home production into account. The utility measure requires the model, since utility is based on the calibrated preference parameters for each of the goods. The two measures are illustrated in figure 4, which shows their cross-sectional distribution. I focus on the variance as a measure of inequality since it facilitates decompositions and has often been used in the literature on consumption inequality (e.g. Lise and Seitz (2011)).

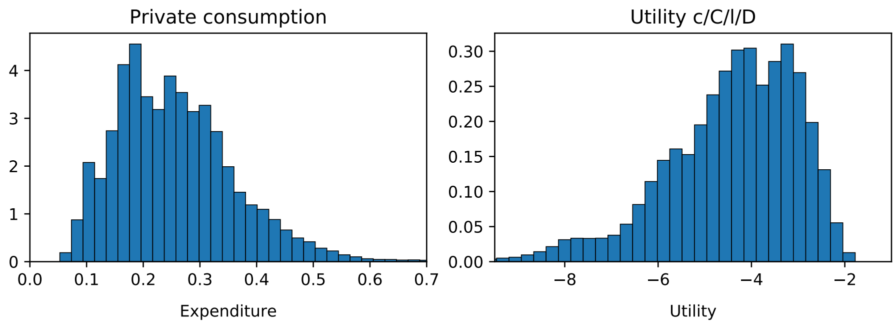

FIGURE 4

Notes: The figure shows the cross-sectional distributions of private consumption and the utility from private/public consumption, leisure and home production for working-age individuals based on the calibrated model.

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. The variances within and between couples are defined as follows (h is the household index):

$$\begin{aligned} \text{Var}[X_i| \text{ within married}] &= E[\text{Var}(X_i|i \in h)] \\ \text{Var}[X_i| \text{ between married}] &= \text{Var}[E(X_i|i \in h)] \end{aligned}$$

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. This leads to the following decomposition for the population variance:

$$\begin{aligned} \text{Var}[X_i] &= p_m \text{Var}[X_i| \text{ within married}] \\ &+ p_m \text{Var}[X_i| \text{ between married}] \\ &+ p_s \text{Var}[X_i| \text{ singles}] \\ &+ \text{Var}[X_i| \text{ between married and single}] \end{aligned}$$

p_m is the share of married individuals in the population and $p_s = 1 - p_m$ the share of singles.

Var$[X_i|$ singles$]$ refers to the variance across single individuals. In addition, there is another component that measures how different the means of singles and married individuals are, which is refered to as the variance between these two groups[23]:

$$\begin{aligned} \text{Var}[X_i| \text{ between married and single}] &= p_s \cdot (E(X_i| \text{ single}) - E(X_i))^2 \\ &+ p_m \cdot (E(X_i| \text{ married}) - E(X_i))^2 \end{aligned}$$

This term captures economies of scale - in terms of consumption, married individuals are able to afford more - and selection into marriage, since the type compositions of the groups of married and single individuals differ. The variances get weighted by the share of the population that is married or single (p_s and p_m).

Table 4 shows the results of the decomposition based on the calibrated model. The first colum reports the decomposition of log private consumption. The single component accounts for 30.3% 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 20.1% of inequality in private consumption, while the between household component accounts for about 49.5%. The variance between married couples and singles is small (0.1%), indicating that the difference in the means of the two groups is small.

TABLE 4: Variance decomposition

	$\log(\mathbf{c_i})$	$\mathbf{u}(\mathbf{c_i}, \mathbf{C_i}, \mathbf{l_i}, \mathbf{D_i})$
Singles (%)	30.3	33.1
Within couples (%)	20.1	5.4
Between couples (%)	49.5	57.3
Betw. sin. and mar (%)	0.1	4.2

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

The second column shows the decomposition for the per-period utility. This measure summarizes inequality in private and public consumption, leisure and home production. In particular, bargaining influences the allocation of leisure in addition to the private consumption good, so that jointly considering these goods is a better indicator of the extent of intra-household inequality. At the same time, the measure takes into account that home production and a part of consumption

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

$$\text{Var}[X_i] = p_m \text{Var}[X_i| \text{ married}] + p_s \text{Var}[X_i| \text{ singles}] + \text{Var}[X_i| \text{ between married and single}]$$

are public goods, and weights each component by the calibrated preference parameter. For the utility measure, the single component accounts for 33% of the total variance, which is a bit larger than for private consumption. The within-couple components accounts for around 5% of the total variance and the between-couple component for 57% of the total variance. Interestingly, the variance between singles and married individuals accounts for about 4% of the variance in the utility measure. This term captures how different the means of utility of married and single individuals are.

For comparison and robustness, I also conducted the decompositions for alternative measures, including a 'full expenditure' measure that weights leisure and home production time by their market value (as in e.g. Lise and Yamada (2017)) and an equivalence scale that compares utility levels by converting them into a money amount. These cases are discussed in more detail in section 5.5. The results from the alternative measures are comparable to the utility measure.

5.2 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

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

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 funtion to income tax rates.[24] 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.[25]

The hypothetical reform increases the progressivity parameter by 0.06. To illustrate the magnitude of the tax change, figure 5 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 $6 - 8\%$, while it increases the average tax rate for high income levels (around 200.000 €) by $\approx 4\%$. 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 (2014). The focus of the policy analysis is mainly to compare the relative magnitude of the different components of the variance decomposition. I also also

[24] I 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.

[25] Computing 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.

conducted the experiment for different changes in the progressivity parameter and the relative importance of the components is similar.

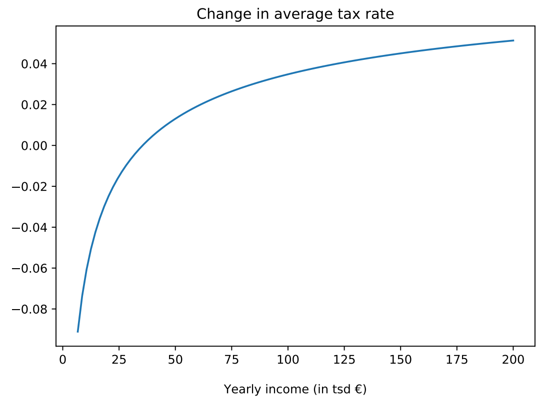

FIGURE 5

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

5.3 Effects on 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. The goal is to express the total percentage change in the variance ($\hat{V} = \frac{V_1 - V_0}{V_0}$) as a weighted sum of the changes of the variance within and between couple, of singles and between couples and singles. This leads to the following decomposition of the change:

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

$\hat{V}_S$ is the percentage change in the variance of singles, $\hat{V}_{M,B}$ and $\hat{V}_{M,W}$ are the changes of the variances between and within married couples and $\hat{C}^{BMS}$ is based on the variance between couples and singles. ω_j is the weight of each growth rate. I separate the effect of the change in the share of the population that is married ($\hat{p}_m$), which introduces a residual, that is typically very small in practice ($\hat{R}$). The terms are described in more detail in appendix C.[26]

The interpretation of the formula is that each component indicates how much it alone contributes (in percentage terms) to the total percentage change in the variance. Consider for example the component for singles ($\omega_1^S \hat{V}^S$). Its contribution to the total change in variance depends on how much the variance of singles changes ($\hat{V}^S$) and on how large the variance of singles was relative to the other variance before the reform (ω_1). The decomposition contains a separate term for the

[26]The decomposition of the change uses the usual calculation rules for growth rates, i.e. the formulas $x \hat{+} y = \hat{x}\frac{x}{x+y} + \hat{y}\frac{y}{x+y}$ and $\hat{xy} = \hat{x} + \hat{y} + \hat{x}\hat{y}$.

change in the share of married individuals in the population ($\hat{p}_m$). The interpretation of this term is that changes in the share of the population that is married contribute to changes in inequality as long as the groups of singles and married individuals differ in their variance. In addition, the fraction of married individuals also enters the term for the variance between singles and couples.

Table 5 shows the results of the decomposition for private consumption and the utility measure. The first row shows the total reduction in inequality. The other rows contain the contribution of each component as a fraction of the total change. The table shows that the variance of log private consumption decreases by 12.7% and e.g. 23.6% of this reduction are due to the single component. The change in the within-couple component accounts for about 16.5% of the total change. For the utility measure, the total decline is by 8.7% and the within component accounts for 6.9% of the total change in variance. Thus, the policy change reduces inequality in terms of private consumption more strongly than for the utility measure. For the utility measure, the component between single and married individuals and the change in the single probability contribute to the total effect, although these effects are relatively small (3.4% in total).

TABLE 5: Variance decomposition - Policy change

	$\log(\mathbf{c_i})$	$\mathbf{u}(\mathbf{c_i}, \mathbf{C_i}, \mathbf{l_i}, \mathbf{D_i})$
Total (%)	-12.7	-8.7
Singles (% of tot.)	23.6	37.9
Within couples (% of tot.)	16.5	6.9
Between couples (% of tot.)	59.8	54.0
Betw. sin. and mar. (% of tot.)	-0.0	2.3
Single probability (% of tot.)	-0.0	1.1
Residual (% of tot.)	-0.0	0.0

Notes: This table shows the effect of the policy change on the variance of private consumption and per-period utility. 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.

The decompositions from tables 5 and 10 focused on the contributions of each component to the total change in variance. This depended both on how reactive a component is to the tax system and how large the pre-reform level of that component is relative to the others. To focus only on the relative *reactiveness* of the between- and within household component, table 6 shows the unweighted variances (e.g. $\hat{V}^{M,B}$ instead of $\omega_2\hat{V}^{M,B}$) and how they change due to the reform. This addresses the question which of the two components reacts more strongly to the reform. The table shows the results for private consumption and for the utility measure. In the first row of sections (a) and (b), the values of the within- and between- couple components are shown. The final column contains the relative size of the two. The second row then shows the percentage change of each variance and the relatize size of the changes.

TABLE 6: Unweighted change in variance

	Var. within couples	Var. between couples	Within rel. to between
(a) $\log(\mathbf{c_i})$			
Variance ($\psi_2 = 0.15$)	0.055	0.135	0.407
Change ($\psi_2 = 0.21$)	-0.107	-0.154	0.695
(b) $\mathbf{u}(\mathbf{c_i}, \mathbf{C_i}, \mathbf{l_i}, \mathbf{D_i})$			
Variance ($\psi_2 = 0.15$)	0.163	1.694	0.096
Change ($\psi_2 = 0.21$)	-0.088	-0.077	1.143

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.

For private consumption, the within variance has about 40% the size of the between-component and declines by 10.7% due to the reform, whereas the between-couple variance declines by 15.4%. For the utility measure, the within-variance declines by 8.8% and the between variance by 7.7%.

Taken together, the results from this section show how progressivity affects inequality within and between couples, among singles and between singles and married individuals. For private consumption, the contribution of the within couple component to the total change of the reform is 16.5%. For the utility measure, it contributes 6.9%.

5.4 Marriage outcomes

In this section, I discuss how marriage outcomes are affected by the tax reform, focusing on differences by ability and gender. In terms of ability, I first focus on differences between individuals *within* the ability groups. Recall that the meeting structure is such that individuals from the same ability group (e.g. types 1, 2 and 3) are more likely to meet each other. The three types from each group will be refered to as the below-average, the average and the above average type. The difference in full-time earnings between these types is approximately 33%, so that the below average type earns 44% of the above average type. The increase in progressivity might improve marriage prospects for the below-average types, since it increases their earnings relative to others. Table 7 investigates this issue in more detail.

The table first shows the probability of getting married (conditional on a meeting) and the Pareto weight and the mean value of match quality for those meetings that result in a marriage. The table reports the mean over the below-average, average and above-average types across the type groups.[27] To focus on the early phase of matching, the table is based on individuals of age 40 or below. The first row of the table shows the share of meetings that result in marriage. For men, the type below the group average is least likely to get married (31%), while the type above the group average has a probability of getting married of 37%.

The mechanism that leads to the differences across types is that meetings are random and the below-average type has a lower income relative to the men that women are likely to meet in the next period. As a result, women prefer to wait for meeting someone else in some cases. This also depends on how high the match quality draw is, since a high match quality make the economic characteristics less relevant. At the same time, individuals with lower earnings are also the least selective and can partially compensate for their lower earnings by accepting a lower bargaining weight. As a result, it is ex-ante not clear if a lower wage translates into a lower marriage probability. The second row of the table shows the bargaining weights, which vary significantly with the type: the below-average type of man is assigned only a Pareto weight of 0.4, whereas the above-average type gets a weight of 0.62.[28] The third row reports the value of match quality. This is *higher* for the below-average type, which reflects that it requires a higher match quality draw for this type to get married.

Interestingly, the picture is quite different for women. The below-average type of women is *more* likely to get married than the above-average type and the difference between these types is less pronounced than for men. Similarly to men, a higher type results in a higher Pareto weight. For women, the difference in match quality between the types is smaller. Men and women differ due to the gender wage gap and the difference in home productivity. Since women are more likely to reduce their labor supply in marriage, their ability type matters less from an economic point of view than the ability type of men.

Section (b) of table 7 shows how these outcomes are affected by the increase in progressivity. The marriage probabilities increase for almost all types. The increase is strongest for below-average men (1.9%). Their mean Pareto weight remains unchanged and the average match quality decreases, since the match quality 'threshold' declines. Below-average females also experience an increase in their marriage probability. For them, the change in the marriage probability is smaller than for men (0.3% vs 1.9%) and their Pareto weight increases. For the other two groups of men, the mean marriage probability increases, while their Pareto weight decreases, which indicates that they get married to 'better' partners than before.

Since the gender difference is relatively strong under status quo, I considered another experiment to see if a stronger increase in progressivity could mitigate the differences in marriage rates more substantially. In this experiment, the progressivity parameter is increased to 0.33, which

[27] For example, the first column is the average for types 1, 4, 7 and 10, which are the lowest in their groups.

[28] The weight reported for males is $1 - \lambda$.

TABLE 7: Marriage market outcomes - within ability group

	M - 1	M - 2	M - 3	F - 1	F - 2	F - 3
(a) Status quo						
Marriage probability	0.31	0.35	0.37	0.36	0.34	0.32
Pareto weight	0.4	0.53	0.62	0.37	0.5	0.59
Match quality	1.12	1.07	1.04	1.06	1.08	1.1
(b) Policy experiment						
Change in mar. (%)	1.9	0.9	0.8	0.3	1.7	1.5
Change in Pareto weight (%)	0	-0.8	-1.5	2.4	1	0
Change in match quality (%)	-0.5	-0.6	0.3	-0.1	-0.6	-0.2

Notes: The table reports the marriage market outcomes under the status quo and after the reform. The columns show the outcomes separately for the below- average, the average and the above-average types of the ability groups.

increases progressivity by 120%. This experiment is illustrated in figure 6. The larger increase makes the increase in marriage rates more substantial; the share of below-average men that is currently married increases by around 5 percentage points between age 23 and 38. Finally, the figure also shows the change in divorce rates for the 40%-increase in progressivity. The figure shows the change in share of individuals that ever divorce (divided by the share ever married) from an ex-ante perspective separately by ability *group*. The change is largest for the highest two ability groups, who become poorer due to the reform and value the economic benefits of marriage more. Interestingly, low types are also slightly *less* likely to divorce. The income effect goes into the other direction for them.

For comparison, table 8 shows marriage rates across ability *groups*, rather than focusing on differences within groups, for the status quo. When comparing ability groups, the lower groups marry at *higher* rates than higher groups. This results from the calibrated parameter for within-group meetings being relatively high. Individual from the lowest ability group have a high probability of meeting someone else from the lowest ability group. They are less picky about match quality, compared to individuals from higher ability groups, since the economic gains from marriage are larger for them. For example, men from the lowest group have a mean match quality of 1.02, while it is 1.11 for the highest ability group. The lower wage, across groups, is therefore reflected in a lower match quality in marriage, rather than a lower marriage rate.

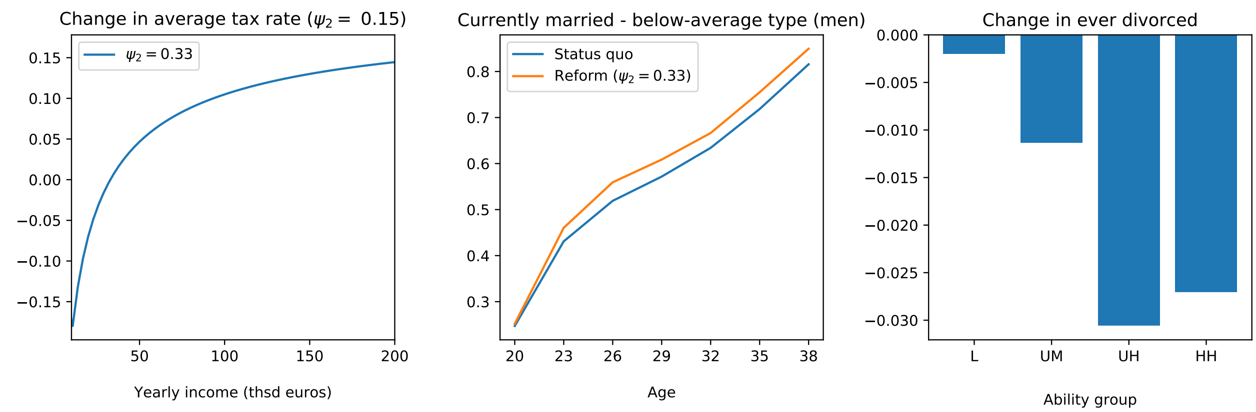

FIGURE 6: Additional results

Notes: The first two figures show the impact of a stronger increase in progressivity on the share currently married of the below-average male type. The left graph illustrates the change in average tax rates relative to the status quo. The third figure shows the change in divorce for the 40%-increase in progressivity by ability group.

TABLE 8: Marriage market outcomes - across ability groups

	M - 1	M - 2	M - 3	M - 4	F - 1	F - 2	F - 3	F - 4
Marriage probability	0.37	0.34	0.33	0.32	0.42	0.35	0.31	0.29
Pareto weight	0.46	0.51	0.54	0.56	0.42	0.48	0.52	0.55
Match quality	1.02	1.08	1.10	1.11	0.98	1.08	1.12	1.16

Notes: The table shows the mean of the marriage market outcomes for each ability group.

5.5 Alternative welfare measures and robustness

For further comparison and robustness, table 9 reports the results of the variance decomposition for three additional welfare measures. The first measure focuses on consumption inequality and sums the value of private and public consumption for each individual. The other two measures are alternative ways of summarizing the value of expenditure and time use, as opposed to considering inequality in the per-period utility. I consider 'full consumption', as in e.g. Lise and Yamada (2017), and an equivalence scale that converts utility to a monetary index.

For private and public consumption, the resulting share of the within-component is relatively small (1.1%). This reflects public consumption being a relatively large part of total expenditure. The share is lower than the within-share found in Lise and Seitz (2011), partly because the fraction of public consumption is higher here. Interestingly, the component between singles and couples accounts for 13.3% of the variance of total consumption, indicating that married individuals are on average richer in terms of total consumption. For private consumption only, this component was small (0.1%). Following Lise and Yamada (2017) 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. For full expenditure, there still is a mean difference between married and single individuals. The within-couple component is somewhat smaller than for the utility measure and accounts for 3.8% of the total variance.

The third measure is an equivalence scale (see Pendakur (2018) for an overview). The measure is also based on comparing static utility levels across people and converts them into a monetary amount for a reference person. 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 chosing leisure and home time. For the equivalence scale, the intra-household component accounts for 5.7% of the variance. The main difference to the other measures is that the single component becomes smaller, while the between-couple component is larger.

TABLE 9: Variance decomposition - alternative measures

	$\log(c_i + C_i)$	Full expenditure	Equiv. scale
Singles (%)	33.0	35.4	23.8
Within couples (%)	1.1	3.8	5.7
Between couples (%)	52.6	52.6	65.4
Betw. sin. and mar (%)	13.3	8.2	5.1

Notes: This table computes the variance decomposition for the status quo for the alternative measures and reproduces table 4 for these measures.

Table 9 shows the decomposition of the effect of the policy for the three alternative measures. The reduction in the variance of private and public consumption is quite similar to the reduction in private consumption only. Similarly, the other two measures are quite close to the utility measure in table 5. The equivalence scale assigns the largest share to the within-household component (6.7%). A difference between the equivalence scale and full expenditure is that the sign of the change for the single probability component and the component between singles and married individuals differs.

TABLE 10: Policy change - alternative measures

	$\log(\mathbf{c_i} + \mathbf{C_i})$	Full expenditure	Equiv. scale
Total (%)	-12.1	-10.0	-9.0
Singles (% of tot.)	27.3	30.0	21.1
Within couples (% of tot.)	0.8	4.0	6.7
Between couples (% of tot.)	68.6	65.0	70.0
Betw. sin. and mar. (% of tot.)	0.8	-1.0	3.3
Single probability (% of tot.)	2.5	3.0	-1.1
Residual (% of tot.)	-0.0	-0.0	-0.0

Notes: This table shows the effect of the policy change on the variance of private consumption and per-period utility. 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.

For further robustness, I compared the results to alternative inequality measures instead of the variance. Similar decompostions 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 Life-time inequality

An alternative approach is to focus on life-time inequality and welfare, rather than cross-sectional inequality. The main advantage is that the expected life-time utility of individuals aggregate all effects through which the marriage market influences welfare, whereas the cross-sectional measures are a more partial picture of different dimensions of inequality.

5.6.1 The comparison case

In the context of life-time welfare, it is not possible to conduct decompositions similar to the within-between decomposition from the previous section, since individuals marry and divorce at different times of their life. Thus, I use a different approach to assess the role of the marriage market, based on holding policy functions constant across policy regimes. The interpretation of this experiment is to ask what the effects of the reform would be if marriage and divorce decisions and the bargaining weights were exogenous and would not react to the change in tax policy.

More concretely, consider changing the progressivity parameter from the current level ψ_2^0 to a new level ψ_2^1. To assess the role of the marriage market, I compare the outcomes of the full model to the outcomes of a version of the model where marriage and divorce decisions and bargaining weights are the same as they would have been under the old tax schedule (conditional on state variables). As an example, suppose that there are two potential spouses (f and m) that meet on the marriage market and have the following characteristics of idiosyncratic productivity, ability, assets and kids:

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

In addition, some value of match quality is drawn. Suppose that under the current tax schedule and the calibrated parameters, the marriage decision would be $M \in \{0, 1\}$ and the initial Pareto weight, if marriage takes place, is given by λ_{ini}. Now imagine that the progressivity of the tax schedule changes. In the exogenous marriage market case, whenever two individuals with these characteristics meet, they are exogenously assigned the marriage decision M and the initial Pareto weight λ_{ini}, instead of being able to make these decisions based on the new circumstances. Similarly, the bargaining weight in marriage is updated by exogenously reproducing the mapping between the state variables of married couples and the potentially new bargaining weight. The model is re-solved with this procedure, so that agents form their expectations over the future taking into account that the bargaining weight will be exogenously set. By construction, solving the exogenous marriage market case for the current tax system ($\psi_2^1 = \psi_2^0$) gives identical outcomes to the full model. This is useful, because alternative comparison cases (such as forcing couples to always have a bargaining weight of 0.5) would change all outcomes, in terms of e.g. marriage, divorce and labor supply, for the current tax system.

The exogenous marriage market case is related to the exercise conducted by Knowles (2012), who compares a bargaining model to a version with fixed bargaining weights to contrast the implications of unitary and collective household models. The exogenous marriage market case can be seen as the 'error' one would make with a 'unitary' version of the model in which bargaining weights are exogenously given, and where marriage and divorce are exogenous conditional on state variables.[29]

Note that the exogenous marriage market case has different implications for how intra-household inequality changes with progressivity. When bargaining 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 bargaining weights if it leads to a reduction in their labor supply and home production hours increase less than work hours are increased. The overall effect depends on the home production technology. With flexible bargaining weights, progressivity directly affects relative consumption by changing the bargaining weight.

[29] In practice, most models with exogenous bargaining weights assume a single fixed weight for *all couples* (e.g. 0.5 in Guner, Kaygusuz, and Ventura (2011)), whereas the exogenous bargaining case here assigns one exogenous weight *for each type of couple*. In addition, this weight is allowed to change *exogenously* over time.

5.6.2 Results

Figure 7 starts by showing the overall welfare change for an increase in progressivity for each ability type, averaged over men and women. The welfare change is expressed in terms of equivalent variation. This is the percentage change in consumption in each state that makes the individuals indifferent between the status quo and the reform.[30] The equivalent variation is shown for the full model and for the exogenous marriage market case. The figure indicates that, as expected, low types gain on average from the marriage market adjustment, while high types tend to lose.

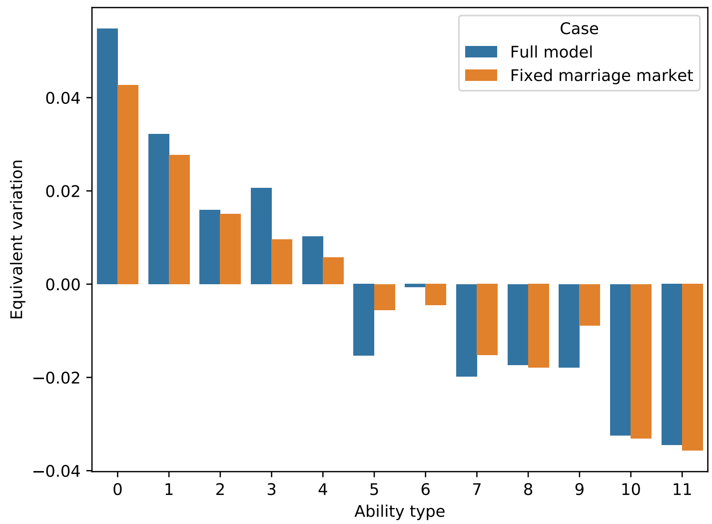

FIGURE 7: Welfare impact along the ability distribution

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

TABLE 11

	1	2	3	4	5	6
Rel. welfare difference	0.22	0.14	0.06	0.54	0.44	0.64

	7	8	9	10	11	12
Rel. welfare difference	5.86	0.23	0.03	0.51	0.02	0.03

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 7.

To illustrate the magnitude of the difference between the two cases, table 11 summarizes the difference in the welfare change between the two cases. The table reports the difference in the

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

equivalent variation between the full model and the case without marriage market adjustments, divided by the welfare change for the full model (in absolute terms). In terms of figure 7, this corresponds to the difference between the bars, divided by the bar for the full model. This is a way to relate to the size of the marriage market adjustments to the total effect of the reform. The numbers are relatively heterogenous along the distribution. For some types, such as the lowest type, for whom the difference has 22% the size of the total effect of the reform, or types 5-6, for whom this value ranges from $44\% - 64\%$, the change is relatively large. For other type, it is rather low, such as for the two highest types, for whom the marriage market effects are only $2-3\%$ of the total effect. In interpreting the differences along the distribution, it is important to keep in mind the meeting structure in the economy, since individuals from each 'type group' are more likely to meet each other. Consider, for example, the case of type 4, which is in the lower middle of the overall income distribution, but the lowest type of its group (consisting of types 4, 5 an 6). As a result, the gain of this type from the marriage market adjustments is larger than the gains from type 3, who is the highest type in their ability group.

5.7 Further experiments

Two important ingredients of the model are the presence of search frictions and the within-ability-group meeting. To study how they affect the outcomes, I conducted further experiments in which I varied the strength of these components. Studying the impact of search frictions on inequality is also interesting because the rise of online dating has arguably led to a large reduction in search frictions. In the long-run, this could affect single rates, assortative matching and intra-household inequality, and ultimately the need for redistribution through a progressive tax system.[31]

In the first experiment, I *increase* the strength of the search friction by making meetings less frequent relative to the benchmark model. The probability of participating in the marriage market is reduced by half. The most direct effect of this change is that there are fewer meetings and the share of singles increases, which increases inequality. The interesting aspect of the experiment is that individuals become less picky about partners and marry in cases where they would not have married if the search friction was lower. As a result, it can also become more likely that richer individuals get married to poorer ones, which can *reduce* inequality between households. At the same time, this effect would *increase* intra-household inequality because couples in which one spouse has a significantly higher earning potential are unlikely to share resources equally. Taken together, the effects of singleness and inequality within- and between couples can go in different directions, which makes it an interesting question how important these channels are relative to each other.

The results from this exercise are summarized in table 12. The rows from the table show the

[31] While some empirical evidence shows that the availability of online dating influences matching outcomes (Hitsch, Hortaçsu, and Ariely (2010), Lee (2016) less is known about how these changes might impact total inequality, once the new technologies are widely used.

decomposition that was previously used, in which the change in the total variance is decomposed into the variance of singles, between singles and married couples and within and between couples, as well as a part due to the change in the single rate. The sum of the components is equal to the value reported in the first column. Focusing on the utility measure of inequality, the table indicates that all components play a role in determining the total impact. The reduction in the between-couple variance (-1%) dominates slightly, leading to an overall decrease in cross-sectional inequality. Interestingly, the within-couple component contributes to *lowering* inequality slightly, but is fairly small (-0.3%). The changes in life-time welfare are shown in A2 in appendix D. All types lose from the increase of the search friction. This loss is partly due to the increased time individuals spend as singles. In addition, the loss also comes from a reduction in the average match quality, since individuals are less able to wait for draws with a high value of the non-economic match quality (θ).

TABLE 12

	Total	Singles (%)	Within couples (%)	Between couples (%)	Betw. sin. mar. (%)	Single prob. (%)
(a) $\log(\mathbf{c_i})$						
Stronger search friction	-0.8	0.4	-0.3	-1	0.2	-0.1
No assortative meeting	-1.4	1.9	11.5	-16.1	1.1	-0.2
(b) $\mathbf{u}(\mathbf{c_i}, \mathbf{C_i}, \mathbf{l_i}, \mathbf{D_i})$						
Stronger search friction	-4.4	-2.2	-0.3	-2.8	0.6	0.4
No assortative meeting	-1.3	3.3	2.4	-13.3	4.9	0.5

Notes: Each row indicates the **change** of the variance (in percent) due to the experiment and the decomposition of the change. For each row, the values from columns indicated as % add up to the total change of that is reported in the first column.

In the second experiment, I study the case where meetings are perfectly random, rather more likely to occur within the ability groups. This varies another dimension of the search friction, since individuals are now more likely to meet potential partners with incomes different from their own.[32] There are three main effects of replacing assortative meetings by random meetings. The first one is that this decreases assortative mating, which leads to less inequality between couples. The effect of assortative mating on between-couple inequality has received a considerable amount of attention in the academic and public discussion (see e.g. Greenwood et al. (2016), Fernandez, Guner, and Knowles (2005)). A less well-studied aspect is to what extent the decline in assortative mating also leads to an increase in intra-household inequality, that results from spouses having

[32] The experiment is also related to some concerns in the literature that the fact that individuals largely meet similar individuals in their community, which is sometimes called 'segregation' (see Fernández (2002) or Wilson (1997)), leads to inequality.

more unequal bargaining positions.[33] The relative importance of these two channels depends on the share of public goods and the calibrated preference parameters for leisure. Finally, there is a third channel, since assortative meeting also affects single rates: when meetings are random, more matches are rejected while individuals try to find their preferred partner. Table 12 shows the results for this experiment. For private consumption, the between component reduces the total variance substantially (-16.1%), while the within component increases inequality by 11.5%. Turning to the utility measure, the off-setting effect of the within-couple component is smaller ($+2.4\%$ vs -13.3%). Still, the total effects deviates substantially from the between-couple component alone, which would decrease inequality by 13.3%, while the eventual decrease is only by 0.7%. Figure A2 also shows the results for life-time welfare. The change is most pronounced for the lowest and highest ability group.

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 effect, 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 20.3% 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 5.4% 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 change in inequality is 16.5% for private consumption and 6.9% for the utility measure. The increase in progressivity increases marriage rates, in particular for men who have a relatively low ability, which further contributes to the reduction in inequality.

[33] Focusing on consumption inequality, Lise and Seitz (2011) show that an increase in assortative mating can explain both an increase in between-HH and a decrease in within-HH inequality.

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.

Blundell, Richard and Ian Preston. 1998. "Consumption inequality and income uncertainty." *The Quarterly Journal of Economics* 113 (2):603–640.

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-André and Costas Meghir. 2015. "Intrahousehold inequality." In *Handbook of Income Distribution*, vol. 2. Elsevier, 1369–1418.

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.

Fernández, Raquel. 2002. "Education, segregation and marital sorting: theory and an application to the UK." *European Economic Review* 46 (6):993–1022.

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.

Gayle, George-Levi and Andrew Shephard. 2018. "Optimal taxation, marriage, home production, and family labor supply." .

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.

Hitsch, Gunter J, Ali Hortaçsu, and Dan Ariely. 2010. "Matching and sorting in online dating." *American Economic Review* 100 (1):130–63.

Holter, Hans A, Dirk Krueger, and Serhiy Stepanchuk. 2014. "How does tax progressivity and household heterogeneity affect Laffer curves?" .

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.

Lee, Soohyung. 2016. "Effect of online dating on assortative mating: Evidence from South Korea." *Journal of Applied Econometrics* 31 (6):1120–1139.

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. 2017. "Household sharing and commitment: Evidence from panel data on individual expenditures and time use." .

Low, Hamish, Costas Meghir, Luigi Pistaferri, and Alessandra Voena. 2017. "Marriage, Labor Supply and the Social Safety Net." .

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." *UCLA CCPR Population Working Papers* .

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. 2017. "The impact of divorce laws on the equilibrium in the marriage market." .

Santaeulàlia-Llopis, Raül and Yu Zheng. 2017. "Why Is Food Consumption Inequality Underestimated? A Story of Vices and Children." .

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

Wilson, William Julius. 1997. *The truly disadvantaged: The inner city, the underclass, and public policy*. University of Chicago Press.

Appendix

Appendix A: Computational details

Solving the model requires finding a fixed point between the distributions of partners that individuals take into account when forming expectations over the future and the implied distributions of singles, which depends on the marriage and divorce choices that individuals make. The basic algorithm is to start with a guess for the distributions of singles, based on simulating the model with singles only, solve the life-cycle problem and update the guess based on the simulating the distributions of singles. This is repeated until convergence.

The life-cycle problem

The life-cycle problem can be solved recursively starting in the terminal period T. 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. Since the temporary wage shocks are i.i.d. and do not influence continuation values, one can solve the value function on a grid for *total resources*, rather than assets, and interpolate values for each combination of temporary shocks. Thus, the temporary shocks do not introduce a separate state variable.

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. During the simulations, potential partners can always be drawn from the pool of currently available singles in the simulation. The guesses for the distributions only enter the simulation through the expected value functions, that determine bargaining weights and marriage and divorce decisions. In the end, the guess for the distributions is updated (with a weight κ on the old distributions).

Implementation

The model implemented with Python and Numba and solved on a computing cluster using MPI. To get good performance out of MPI for this type of model, it is useful to combine it with OpenMP on each node to limit communication time and 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:

$$\begin{aligned} \text{TI} &= \frac{1}{N}\sum_{i}^{N}\frac{X_i}{\bar{X}}\log\left(\frac{X_i}{\bar{X}}\right) \\ \text{MLD} &= \frac{1}{N}\sum_{i}^{N}\log\left(\frac{\bar{X}}{X_i}\right) \end{aligned}$$

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. (%)	Single prob. (%)	Res. (%)
(a) $\mathbf{c_i}$							
MLD ($\psi_2 = 0.15$)	0.1	30	20.2	49.5	0.2	-	-
Change ($\psi_2 = 0.21$)	-12.9	-2.9	-2.1	-7.9	0.3	-0.2	-0
(b) Equivalence scale							
MLD ($\psi_2 = 0.15$)	0.1	22.2	5.5	65.8	6.4	-	-
Change ($\psi_2 = 0.21$)	-9.9	-1.8	-0.6	-7	1	-1.5	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. (%)	Single prob. (%)	Res. (%)
(a) $\mathbf{c_i}$							
Theil index ($\psi_2 = 0.15$)	0.1	29	19.5	51.3	0.2	-	-
Change ($\psi_2 = 0.21$)	-13	-2.7	-2	-8.3	-0.3	0.3	-0.1
(b) Equivalence scale							
Theil index ($\psi_2 = 0.15$)	0.1	17.5	6.3	70.1	6.1	-	-
Change ($\psi_2 = 0.21$)	-10.8	-1.5	-0.8	-8.1	-1.5	1.1	-0.1

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:

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

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: Additional figures

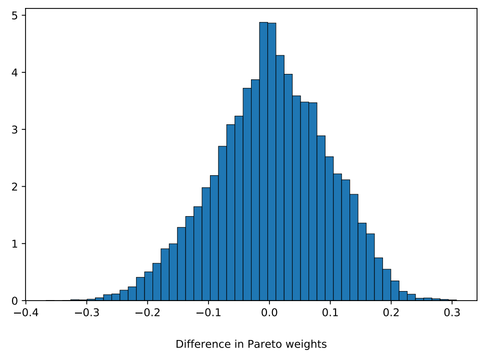

FIGURE A1

Notes: The figure shows the influence of the possibility of future marriage on bargaining weights. For each marriage that occurs in the calibrated model, the bargaining solution was computed for the actual outside option and for the hypothetical case where the outside option is staying single forever. The figure plots the difference between the actual bargaining weight and this alternative weight.

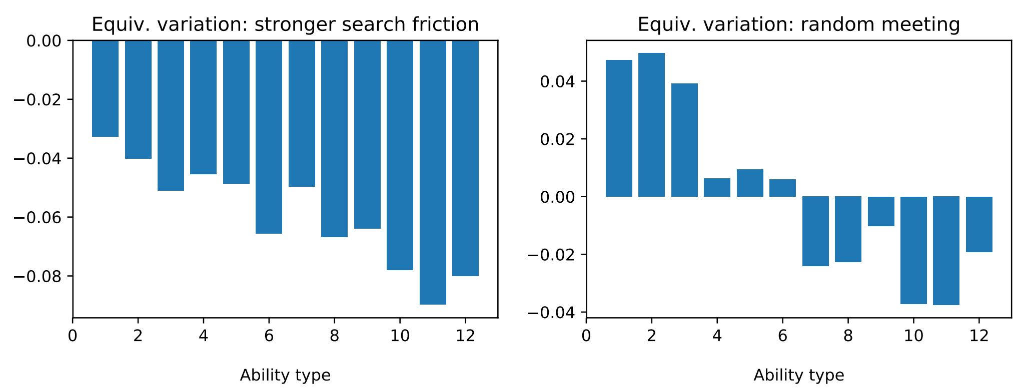

FIGURE A2

Notes: The figure shows welfare changes for the two experiments (in terms of the equialent variation)