NBER WORKING PAPER SERIES

MARRIAGE MARKET AND LABOR MARKET SORTING

Paula A. Calvo
Ilse Lindenlaub
Ana Reynoso

Working Paper 28883
http://www.nber.org/papers/w28883

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2021

We thank participants at the STLAR Conference at St. Louis Fed, the Cowles Macroeconomics Conference at Yale, UPenn, University Washington St. Louis, Frankfurt University, the H2D2 Research Day at University of Michigan, the Search and Matching Workshop at the University of Essex, the Conference on Competitive Equilibrium with Gross Substitutes, with Applications to Problems in Matching, Pricing, and Market Design at Brown University, UT Austin, Universidad de San Andres, University of Michigan, CSWEP Paper Sessions: Gender Topics at the ASSA/AEA 2021 Meetings, Midwest Economics Association 2021 Annual Meeting, University of Maryland, UC San Diego, University of Nottingham, University of Melbourne, University College London, the University of Otago, and the University of Chile for useful comments. The paper also benefitted from discussions with Joseph Altonji, Hector Chade, Limor Golan, and Esteban Peralta. We thank Valentina Goetz for outstanding research assistance. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Marriage Market and Labor Market Sorting
Paula A. Calvo, Ilse Lindenlaub, and Ana Reynoso
NBER Working Paper No. 28883
June 2021
JEL No. D1,J01,J12,J16,J24,J31

ABSTRACT

We build a novel equilibrium model in which households' labor supply choices form the link between sorting on the marriage market and sorting on the labor market. We first show that in theory, the nature of home production – whether partners' hours are complements or substitutes – shapes marriage market sorting, labor market sorting and labor supply choices in equilibrium. We then estimate our model on German data to assess the nature of home production in the data, and find that spouses' home hours are complements. We investigate to what extent complementarity in home hours drives sorting and inequality. We find that the home production complementarity – by strengthening positive marriage sorting and reducing the gender gap in hours and labor sorting – puts significant downward pressure on the gender wage gap and within-household income inequality, but it fuels between-household inequality. Our estimated model sheds new light on the sources of inequality in today's Germany and – by identifying important shifts in home production technology towards more complementarity – on the evolution of inequality over time.

Paula A. Calvo
Department of Economics
Yale University
28 Hillhouse Av.
New Haven, CT 06520
paula.calvo@yale.edu

Ilse Lindenlaub
Department of Economics
Yale University
28 Hillhouse Avenue
New Haven, CT 06511
and NBER
ilse.lindenlaub@yale.edu

Ana Reynoso
Department of Economics
University of Michigan
611 Tappan Avenue
Ann Arbor, MI 48109
areynoso@umich.edu

An online appendix is available at http://www.nber.org/data-appendix/w28883

1 Introduction

Positive assortative matching is a defining feature of both the labor and the marriage markets and has important implications for inequality. On the marriage market, the matching of partners with similar education impacts both within- and between-household income inequality. Moreover, positive sorting in the labor market between workers and firms or jobs reinforces wage inequality across skills. But even though inequality in economic outcomes results from agents interacting in *both* the marriage and the labor market, how the interplay of the two markets *jointly* shapes inequality has not been studied before.

This paper shows that sorting in the marriage and the labor market are linked by households' time allocation choices — how time is divided between market work and home production — and studies how these interconnected markets affect inequality. We build a novel equilibrium model with rich heterogeneity and sorting on both markets and show that, in theory, the nature of the home production technology shapes equilibrium. If spouses' home hours are *complementary*, a 'progressive' equilibrium emerges in which spouses share household tasks and supply similar market hours, and there is positive sorting on both marriage and labor markets. We then estimate our model to investigate the nature of home production in the data. We find that partners' home production time is indeed complementary in today's Germany and this complementarity has become stronger over time. Analyzing inequality shifts, we find that this technological change in home production is a major driver of reduced gender disparities between 1990 and 2016. Importantly, increases in positive assortative matching in *both* the marriage and labor markets further mitigated gender disparities in Germany over the last decades.

Three sets of facts from the German Socioeconomic Panel (henceforth, GSOEP) show a salient relationship between the marriage and the labor market and motivate our analysis. First, as is well documented in the literature, there is positive assortative matching both on spouses' education in the marriage market and between workers' education and jobs' skill requirements in the labor market. Importantly, there is a gender gap in labor market sorting where, conditional on education, men work in more demanding jobs than women. Second, men and women who are more strongly sorted in the marriage market (i.e. those whose education is more similar to their partner's education) are also more strongly sorted in the labor market (i.e. they tend to have the 'right' education level for the jobs they perform). Third, households' labor supply choices form an important link between the two markets: Spouses with more similar education tend to split their time similarly between market and house work; and conditioning on hours worked, the gender gap in labor market sorting is significantly smaller.

We capture these observed features in a novel equilibrium model in which households' endogenous labor supply choices form the link between the marriage and the labor market. The model is static, and individuals who differ in skills face three decisions. First, in the marriage market, men and women choose whether and whom to marry. Second, each household formed in the marriage stage collectively decides on its members' market and home production hours (where the home hours produce the household's public good), as well as their private consumption. Last, in the labor market, individuals match with firms of different productivity (where we use firms interchangeably with 'jobs'), determining their wages.

The crucial twist of our model is that, in the labor market, firms value *both* workers' skills and hours worked, since hours worked increase individuals' productivity.[1] Matching between workers and firms is then based on workers' *effective* skills—an increasing function of both skills and hours—and firms' productivity. Since the household's time allocation depends on both partners' skills and impacts the jobs they match with on the labor market, marriage market sorting affects labor market sorting. At the same time, when making their marital and household labor supply choices, individuals internalize that an increase in labor hours improves job quality and wages, affecting the value from marriage. Therefore, labor market sorting also affects marriage market sorting. This interrelation between the two markets and sorting margins is the unique feature of our model but also makes the problem complex.

We focus on a tractable transferable utility (TU) representation of our model and characterize two benchmark equilibria depending on the model's primitives. Both feature positive assortative matching between workers and jobs in the labor market driven by productive complementarities. However, the two equilibria differ in household and marriage outcomes depending on properties of the home production function. On the one hand, if home production exhibits complementarity in partners' time inputs, a *monotone equilibrium* arises, characterized by positive sorting in the marriage market and labor hours that are increasing in both own and partner's skills. This equilibrium reflects a 'progressive' economy with a high frequency of two-earner households and where spouses are similar in terms of skills and their split between work and home production. The complementarity in home hours is therefore a force towards positive marriage sorting as well as balanced labor supply, labor market sorting and pay across gender. This leads to a narrow gender wage gap, low within-household income inequality, but high inequality between households. On the other hand, if partners' time inputs are substitutable in home production, a *non-monotone* equilibrium arises, featuring negative assortative matching in the marriage market and labor hours that are increasing in own but *decreasing* in partner's skill. This equilibrium reflects a 'traditional' economy with a high degree of household specialization and disparity in partners' skills—features that widen the gender wage gap and within-household income inequality but put downward pressure on between-household inequality.

The key insight from our model is that marriage and labor market sorting are linked in an intuitive way by households' labor supply choices. The nature of this link depends on whether spouses' hours in home production are complementary or substitutable, a feature that needs to be investigated empirically.

We then study the nature of the home production technology and its role for inequality in the data—both in the cross-section and over time. To do so, we minimally augment our baseline model to capture additional sources of observed heterogeneity while preserving its parsimony and key mechanism. First, we introduce three shocks: Marriage taste shocks to allow for mismatch in the marriage market; labor supply shocks to capture time use heterogeneity within each couple-type; and a random component

[1] We base this assumption on our own evidence of a positive impact of hours worked in the labor market on *hourly* wages in the GSOEP, both in Figure O3 and Table 6, column (3); and also on empirical evidence from the literature arguing that more hours worked lead to higher productivity, especially if it is costly to hand off clients, patients or customers to the next worker on the shift, for instance due to increased coordination costs (e.g. Goldin, 2014).

to workers' skills to account for mismatch in the labor market. Second, we further parameterize our model allowing for gender differences in both home and labor market productivity (the latter can also be interpreted as discrimination) that will be disciplined by the data. We show that this model is identified.

We first estimate our model on data from MODERN GERMANY—our benchmark estimation, which focuses on West Germany from 2010 to 2016—and find that spouses' home production hours are complementary. Our model matches key targeted features of the marriage market equilibrium (such as the degree of marital sorting and the high correlation of home hours within couples) and the labor market equilibrium (such as moments of the wage distributions). To further validate the model, we show that it also reproduces key features of the equilibrium that were *not* targeted in estimation: the three stylized facts outlined above as well as our measures of household and gender wage inequality.

In order to showcase our model's mechanism, we conduct comparative statics of the gender wage gap and within/between household wage inequality with respect to three parameters that significantly impact inequality: (i) complementarity of partners' home production time, (ii) women's relative productivity in home production, and (iii) women's relative productivity in the labor market. Our insights are the following: First, eliminating gender asymmetries in productivity (whether at home or at work) naturally reduces the gender wage gap. But, interestingly, an increase in complementarity of partners' home production hours has qualitatively similar effects. Second, a decline in the gender wage gap goes hand in hand with a decline of gender gaps in labor hours and labor market sorting, and with an increase in marriage market sorting. Third, the gender wage gap moves hand in hand with *within*-household inequality but in opposite direction as *between*-household inequality.

Having well understood our model mechanism, we then focus on GERMANY OVER TIME and investigate how our model rationalizes the large decline in gender and within-household income inequality and the increase in between-household inequality between 1990 and 2016. To this end, we re-estimate our model on West German data from the 1990s and then compare it to our baseline estimation. The model estimates reveal significant changes in home production over time with today's Germany being characterized by stronger complementarity in spouses' home hours and increased relative productivity of men, indicating a switch towards a more 'progressive' economy (the monotone equilibrium of our model). These changes in home production technology account for around 70% of the observed decline in the gender gap and for the entire drop in within-household inequality. In contrast, changes in labor market technology—which we interpret as skill-biased technical change—had very different effects: They *fueled* gender and household inequality across the board, preventing gender gaps from falling even further.

Finally, we find that changes in both marriage market sorting and labor market sorting—which increased by 10% and 8%, respectively—significantly affected these inequality shifts. If sorting patterns had stayed constant at their 1990-levels, gender inequalities would have been wider today and between-household inequality narrower. Intuitively, stronger marriage market sorting over time generated more gender-balanced labor market outcomes—in hours, sorting, and pay. In turn, the increase in labor sorting over the past decades also significantly narrowed gender disparities as it was predominantly driven by women's improved labor sorting, helping them to catch up with men's pay.

2 The Literature

This paper relates to four strands of literature, and we clarify our contribution to each of them.

Gender Gaps in Labor Supply and Pay. A growing literature studies the link between the gender gap in labor supply and the gender gap in pay. The standard channel works through earnings, where family and fertility choices have a permanent effect on the gender earnings gap (Adda, Dustmann, and Stevens, 2017; Dias, Joyce, and Parodi, 2018; Angelov, Johansson, and Lindahl, 2016; Kleven, Landais, and Søgaard, 2019). Because the wage rate is kept fixed in these papers, any gender gap in pay can only be attributed to *earnings* not to *hourly wages* (which is what we focus on). In assuming that hours worked affect workers' productivity in the market, we follow more closely the literature documenting significant labor market returns to hours (Aaronson and French, 2004; Gicheva, 2013; Goldin, 2014; Cortés and Pan, 2019; Bick, Blandin, and Rogerson, 2020).[2] Other work links gender pay gaps to gender differences in preferences for work flexibility (Bertrand, Goldin, and Katz, 2010, Mas and Pallais, 2017, Cubas, Juhn, and Silos, 2020) and to sorting into occupations that require different time inputs (Erosa, Fuster, Kambourov, and Rogerson, 2017). Finally, there is work highlighting the importance of information frictions for gender pay gaps (without considering the marriage market): If employers believe that women have less market attachment relative to men, they get paid less (Albanesi and Olivetti, 2009, Gayle and Golan, 2011).

Our paper builds on this work in that we also propose the gender gap in hours as a key factor behind the gender pay gap. However, in contrast to both the purely empirical and the structural papers we cited, our work takes into account an endogenous marriage market which shapes labor supply choices.

Marriage Market Sorting. A large literature measures marriage sorting in the data and finds evidence of positive assortative matching on education in different countries, and increases in marriage sorting over time (Browning, Chiappori, and Weiss, 2014; Eika, Mogstad, and Zafar, 2019; Greenwood, Guner, Kocharkov, and Santos, 2016; Greenwood, Guner, and Vandenbroucke, 2017). We confirm these findings on positive marriage sorting on education for Germany.

Another approach has been to study marriage market sorting using structural models. Several papers have investigated how pre-marital investments in education interact with marriage patterns in a static framework (Chiappori, Iyigun, and Weiss, 2009 and Fernández, Guner, and Knowles, 2005) or over the life cycle (Chiappori, Costa-Dias, and Meghir, 2018) and how post-marital investments in a partner's career interact with family formation and dissolution (Reynoso, 2019). Further, there is structural work analyzing how exogenous changes in wages, education and family values (Goussé, Jacquemet, and Robin, 2017a), exogenous wage inequality shifts (Goussé, Jacquemet, and Robin, 2017b), the adoption of

[2] Bick et al. (2020) find that hourly wages of U.S. men are non-monotone, increasing until 50 hours per week, and then decreasing. Note that in our sample, hardly anyone (<0.3%) works more than 50 hours per week, justifying that we do not allow for non-monotone effects of hours on wages in our model (we do allow for non-linear effects).

unilateral divorce (Reynoso, 2019), or different tax systems (Gayle and Shephard, 2019) affect household behavior and marriage sorting.[3] Finally, in a household model with *exogenous* marriage sorting (and exogenous labor market), Lise and Seitz (2011) analyze the effect of an increase in marriage sorting on between and within household consumption inequality.

Like in these papers, marriage market sorting is an important margin in our model. While education is exogenous in our setting, we could think of the choice of how many hours to work as an 'investment' in individuals' effective skills. But this investment happens *post-marriage* market and *pre-labor* market, and therefore is impacted by marriage sorting while impacting labor market sorting, so the timing is different than in existing work. Crucially none of these papers endogenizes the labor market or features labor market sorting, which is the key addition of our work.[4]

Labor Market Sorting. A body of literature investigates sorting on the labor market, documenting positive assortative matching between workers and firms (Card, Heining, and Kline, 2013; Hagedorn, Law, and Manovskii, 2017; Bagger and Lentz, 2018; Bonhomme, Lamadon, and Manresa, 2019); or workers and jobs (Lindenlaub, 2017; Lise and Postel-Vinay, Forthcoming; Lindenlaub and Postel-Vinay, 2020) without taking the marriage market into account. Our paper is perhaps closest to Pilossoph and Wee (2019b) who consider spousal joint search on the labor market to explain the marital premium, but taking marriage market sorting as given. Our contribution is to explore how the forces that determine who marries whom shape labor market sorting and pay.

Interplay between Marriage and Labor Market. Our work is most related to a nascent literature on the interplay between marriage and labor markets. This research has focused on the effects of spouses' joint labor search (Pilossoph and Wee, 2019a and Flabbi, Flinn, and Salazar-Saenz, 2020), changes in wage structure and home technology (Greenwood, Guner, Kocharkov, and Santos, 2016, Chiappori, Salanié, and Weiss, 2017), and changes in the skill premium (Fernández, Guner, and Knowles, 2005) on marital sorting and household inequality keeping the labor market in *partial* equilibrium.

To the best of our knowledge, this is the first paper that features both the marriage market and the labor market in *equilibrium* with market clearing, price determination and sorting in both markets. Jointly considering marriage and labor market *sorting* is novel and so is our mechanism of how the two sorting margins are linked (i.e. through endogenous labor supply).

3 Descriptive Evidence

3.1 Data

We use two different data sources: The German Socioeconomic Panel (henceforth GSOEP) is a yearly household survey of around 25,000 individuals (including the surveyed households' head and the spouse).

[3]In an influential paper, Voena (2015) also focuses on the adoption of unilateral divorce and its effects on household behavior, especially asset accumulation. In her paper, the marriage market is exogenous.

[4]The exception is Fernández et al. (2005) who endogenize the wages of their two worker types, low and high skill, but as in the other papers, their model does not feature any labor market sorting.

It contains detailed information on labor market outcomes and time use. We focus on West Germany, 2010-2016. The Employment Survey of 2012 (henceforth BIBB Survey) contains detailed occupational characteristics. Details on the GSOEP are in Online Appendix OC and on the BiBB in Appendix C.3.

3.2 Empirical Evidence

We first present evidence related to sorting in the marriage market, sorting in the labor market, and the interaction between the two. We then highlight that the allocation of hours between labor market and home production is an important link between both markets. A description of the sample restrictions and the construction of the main variables can be found in Online Appendix OC.1.

Marriage Market Sorting. We find evidence of positive assortative matching (PAM) in education in the German marriage market, in line with existing evidence (Eika et al., 2019 for the US and Germany and Greenwood et al., 2016, Greenwood et al., 2017 for the US). Table 1 reports the matching frequencies by education for the period 2010-2016, suggesting that almost 60% of individuals marry someone of the same level of education. The correlation between the education level of spouses, which is our main measure of marriage market sorting, is equal to 0.47.[5] Furthermore, marriage market sorting increased over time. For the period 1990-1996, the correlation between partners' education was 0.44.[6]

Table 1: Marriage Matching Frequencies by Education

	Low Education Men	Medium Education Men	High Education Men
Low Education Women	**0.16**	0.08	0.03
Medium Education Women	0.13	**0.25**	0.11
High Education Women	0.03	0.05	**0.16**

Notes: Low education includes high school and vocational education with less than 11 years of schooling. Medium Education is vocational education with more than 11 years of schooling. High Education is defined as college or more. We consider the maximum level of education attained by each individual and keep only one observation per couple.

Labor Market Sorting. We also document positive assortative matching in the labor market. We do not have firm identifiers in the GSOEP, so we measure labor sorting based on the relationship between workers' and *jobs'* attributes, where a job is defined by the occupation of the individual. The match-relevant attribute of workers in the labor market is 'years of education'. In turn, for jobs we use information on the task requirement of each occupation to construct a measure of its *task complexity* (see Appendix C.3). The correlation between years of education of workers and task complexity of jobs is 0.62, indicating positive assortative matching between workers and jobs on the labor market.

Figure 1 (left) plots the fitted matching function (job attribute as a function of worker characteristic) by gender, conditional on employment (solid lines). Both men and women are positively sorted in the labor market, indicated by a positive slope of the matching function. However, men

[5] As discussed in Chiappori, Dias, and Meghir (2020), *correlation* is one of the measures of sorting that has two desirable properties such a measure should have: a 'monotonicity' condition and whether it captures the case of 'perfectly assortative matching'. Eika et al. (2019) propose an alternative measure of marriage sorting based on the frequency of couples with similar education relative to random matching. This measure equals 1.73 in our sample: Individuals are 73% more likely to marry someone of the same education, relative to random matching.

[6] For years 1990-1996, the sorting measure from Eika et al. (2019) is 1.59, also suggesting increased sorting over time.

Figure 1: Labor Market Matching Function (left); Labor Market and Marriage Market Sorting (right)

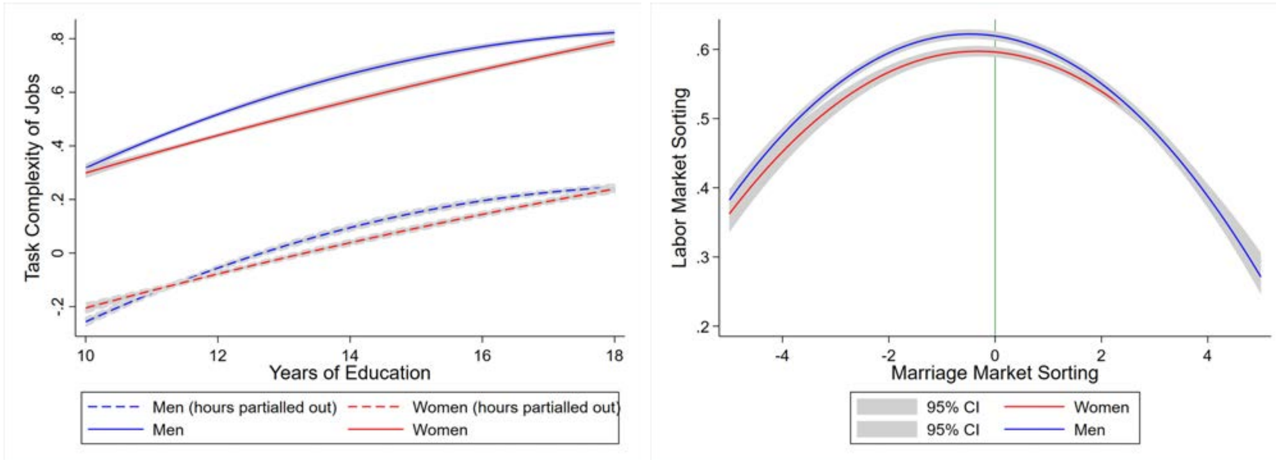

are 'better' matched: for a given education level, men are on average matched to more demanding jobs. This pattern is also reflected in the correlation of worker and job attributes by gender, which is 0.64 for men and 0.62 for women.[7]

Labor Market Sorting and Marriage Market Sorting. Next we assess the relationship between labor market and marriage market sorting. To do so graphically, we measure marriage market sorting by the difference between the years of education of an individual and the years of education of their partner, with 'zero' indicating the maximum amount of sorting. We measure labor market sorting as before as the correlation between workers' years of education and the task complexity of the occupation. We then plot the relation between labor market and marriage market sorting by gender in Figure 1 (right), where the green vertical line indicates maximum marriage market sorting. The striking—and we believe novel—feature is that labor market sorting is maximized when marriage market sorting is maximized, both for men and for women. In Online Appendix OA.1, we substantiate this finding using regressions that control for important covariates.

The Role of Hours. We now provide evidence on a salient link between the two markets: hours worked on the labor market vs. hours spent in home production. First, we document that the time allocation choice is 'impacted' by the partnership status as well as marriage market sorting. Second, we document that at the same time, the time allocation choice is linked to labor market sorting.

As is well documented (Gayle and Shephard, 2019; Goussé et al., 2017b), an individual's time allocation between 'work', 'home production' and 'leisure' is related to their partnership status. While among singles (Online Appendix OA.4, Figure O2, left panel) gender differences in time allocation choices are small, for couples gender differences are pronounced (right panel). Indeed, in couples, women spend about 12.5 hours less per week working on the labor market but about 20 hours more

[7] Differences in labor market sorting across gender are even larger (with a correlation 0.58 for men vs. 0.48 for women) if we do not condition on participation and we treat unmatched individuals as matched to a job with attribute zero. This suggests that non-participation is one of the dimensions through which women are worse matched in the labor market.

in home production compared to their male partners. Neither for couples nor for singles are there significant gender differences in leisure, justifying that we abstract from it in our model below.

We also document the relationship between hours and marriage market sorting. Figure 2 shows the correlation between home production hours within couples (left) and the correlation of labor hours within couples (right), both against our summary measure of marriage market sorting (difference in partners' years of education). Interestingly, both home production hours and labor market hours are more complementary among those partners who are well sorted in the marriage market, as indicated by the inverse U-shape of the hours' correlation function. Note that the pattern for home production is even more pronounced than for market hours, with a stronger positive correlation of home hours among partners with the same education compared to partners with differences in education.[8]

Figure 2: Time Allocation and Marriage Sorting

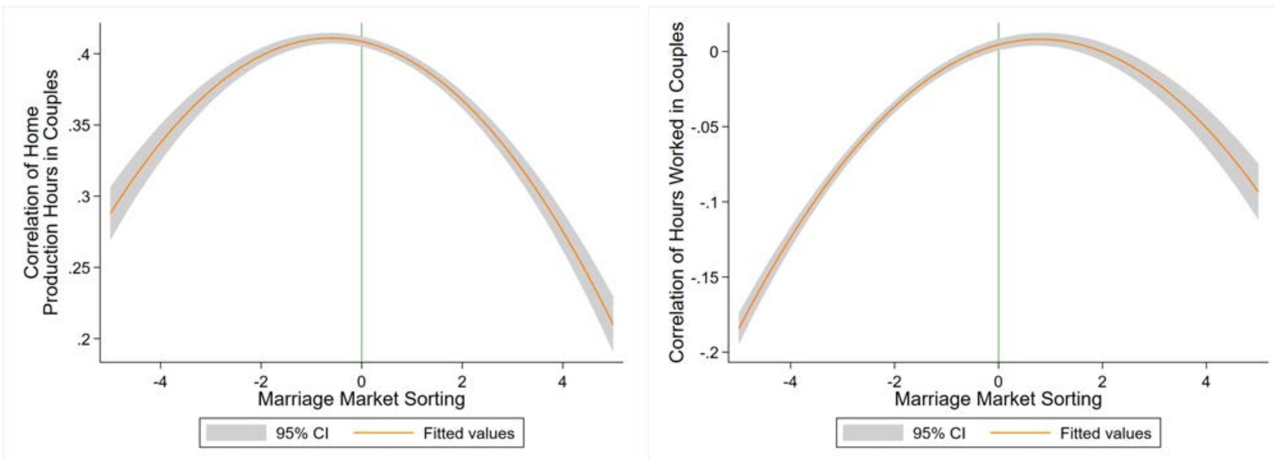

We further explore what drives the complementarity in home production time of spouses by looking at different components of home production, especially since the literature on Family Economics emphasizes household specialization. We find that complementarities are strongest in childcare, and are least pronounced when it comes to housework (see Figure O1 in Online Appendix OA.2).

One concern is that the relationships between marital sorting and complementarities in hours in Figure 2 are based on marriage market sorting bins that pool individuals from different education groups. Not controlling for education allows for the possibility that hours only depend on own education but do not vary with partner's education if, e.g., low (high) educated workers always put low (high) hours independently of the partner's type. Also, the relationships in Figure 2 cannot be interpreted as causal, since there might be other confounding factors or unobserved heterogeneity driving both partners' choice of hours. We discuss and address some of these concerns in Online Appendix OA.3. There, we control for education and other covariates that might be correlated with hours' choices (Table O2). Moreover, we show that partners' complementarities in labor market hours are stronger when we exploit exogenous variation in childcare availability across states and time to instrument for female labor market hours

[8]For consistency between the right and the left panels of Figure 2, we condition on both partners participating in the labor market. The pattern in the left panel also holds if we do not condition on labor market participation.

(Table O3). Finally, we also show in these regressions that the correlation between partners' hours is larger for those who are better sorted in the marriage market, in line with the descriptive evidence of Figure 2. We provide the details of our identification strategy in Online Appendix OA.3.

The second point we stress is that the time split between labor market and home production also relates to labor market outcomes: We first show that in Germany there is a large hourly wage penalty for working part-time, suggesting that hours are a productive input in the labor market. This is in line with evidence by Aaronson and French (2004), Goldin (2014) and Bick et al. (2020) for the US. Figure O3 (Online Appendix OA.5) shows a sizable part-time penalty, especially for women. While full-time women have a wage penalty of 14.7 percentage points relative to full-time men, when they work part-time the wage penalty increases to 26.6 percentage points. Moreover, while few men work less than full-time (less than 10% of employed men), more than 50% of employed women do so, and are thus particularly affected by the documented wage penalties. These effects of hours on wages cannot be interpreted as causal though. To address selection, we identify the effect of hours on the hourly wage in a panel regression model with individual fixed effects below, where we instrument for hours worked (see Section 7.3.1). We again find a significant wage penalty for not working full-time: An increase from 30 to 40 hours per week raises the hourly wage by around 4%.

Finally, we highlight that the number of hours worked is associated with *sorting* on the labor market. Indeed, when accounting for differences in hours worked across gender, the discrepancy in their matching functions shrinks considerably. This is documented in Figure 1, left panel, where the solid lines represent the matching functions by gender and the dashed lines plot the residualized matching functions, after partialling out hours worked. But even when controlling for the number of hours worked, small differences in labor market sorting across gender persist—which must be accounted for by other factors.

In sum, we highlight three sets of facts. First, there is evidence of PAM both in the labor and the marriage market. But on the labor market, men are 'better' matched than women. Second, there is a strong relation between labor market and marriage market sorting, with labor market sorting being maximized when marriage market sorting is. Third, the split between hours worked in the labor market vs. hours spent in home production is a potentially important link between the two markets. We not only show that time allocation choices depend on marriage market sorting but also that they are themselves associated with labor market sorting. Motivated by these facts we now build a model with endogenous labor *and* marriage markets. We also use these facts to justify several assumptions and to guide our modeling choices regarding the link between labor and marriage markets, where we focus on hours. Finally, we come back to these facts when validating our estimated model below.

4 The Model

We start with an overview: Men and women first decide whom to marry based on their education/skills in a competitive market. Each matched household then optimally chooses private consumption and the time allocation between home production and labor market work, which also pins down the public good consumption. Finally, individuals match with heterogenous firms in a competitive labor market. Figure 3 summarizes these decision stages and the endogenous variables they pin down, which we detail next.

4.1 Environment

There are two types of agents, individuals and firms. There is a measure one of firms. Firms are characterized by productivity $y \in \mathcal{Y} = [\underline{y}, \overline{y}]$, distributed according to a continuously differentiable cdf G, with strictly positive density g. Among the individuals, there is an equal measure of men (denoted by subscript m) and women (denoted by subscript f). The overall measure of individuals is one. Both men and women have exogenously given skills: Denote women's skills by $x_f \in \mathcal{X}_f = [0, \overline{x}_f]$, where x_f is distributed with a continuously differentiable cdf N_f, with density $n_f > 0$. Analogously, men have skills $x_m \in \mathcal{X}_m = [0, \overline{x}_m]$, distributed according to the continuously differentiable cdf N_m with density $n_m > 0$.

Figure 3: The Decision Stages of Individual $i \in \{f, m\}$ of Skill Type x_i

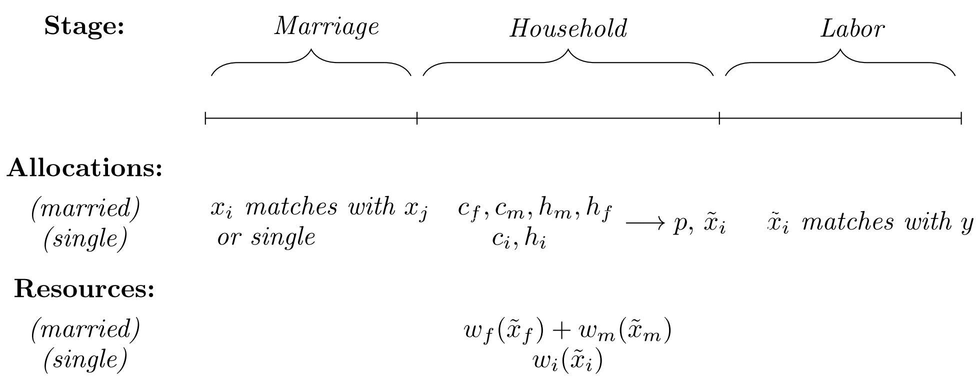

In the marriage market, men and women match on skills, so the relevant distributions for marriage matching are N_m and N_f. In the labor market, however, what matters for output is not only skills but also hours worked, which will be chosen optimally in each couple. Each individual is endowed with one unit of time that can be allocated to paid work in the labor market, denoted by $h_i, i \in \{f, m\}$, or non-paid work at home towards the production of a public good, $\ell_i = 1 - h_i$ (based on Figure O2, which shows no large differences in leisure across gender, we abstract from it). Note that $h_i = 0$ captures non-participation. By increasing hours worked in the labor market, each individual 'invests' in his/her *effective skill* $\tilde{x} := e(x, h)$, $\tilde{x} \in \tilde{\mathcal{X}}$, with endogenous cdf $\tilde{N}(X) := \mathbb{P}[\tilde{x} \leq X] = \frac{1}{2}\mathbb{P}[\tilde{x}_f \leq X] + \frac{1}{2}\mathbb{P}[\tilde{x}_m \leq X]$. We assume that e is twice continuously differentiable, strictly increasing in each argument, (weakly) supermodular.Thus, putting more labor hours is *as if* the worker is more skilled. The effective skill or index $\tilde{x}$ is the output-relevant worker characteristic on the labor market.[9] This assumption that not

[9] We base this assumption on evidence that more hours worked lead to higher productivity and hourly pay (see Aaronson and French, 2004; Gicheva, 2013; Goldin, 2014; Cortés and Pan, 2019; Bick et al., 2020 and our own evidence).

only skills but also hours worked matter for labor market matching means that multiple attributes are matching-relevant even if the actual assignment is simplified and based on the index $\tilde{x}$.

Denote by $z(\tilde{x}, y)$ the output generated by an individual of type $\tilde{x}$ matched to a firm of type y. We assume that production function z is twice continuously differentiable and strictly increasing in each argument. Individuals and firms split the output they generate into wages and profits, where workers use their wages to finance the private consumption good $c_i, i \in \{f, m\}$.

The public good production function is denoted by p, which takes as inputs each couple's hours at home, so that $p(\ell_m, \ell_f)$ is the public good produced by a couple spending $(\ell_m, \ell_f) = (1 - h_m, 1 - h_f)$ in home production (recall hours at home equal the hours *not* spent working).[10] We assume that p is twice continuously differentiable with $p_{\ell_m} > 0, p_{\ell_f} > 0, p_{\ell_m \ell_m} < 0$ and $p_{\ell_f \ell_f} < 0$ and satisfies the Inada conditions $\lim_{h_f \to 0} p_{\ell_f}(1-h_m, 1-h_f) = 0$ and $\lim_{h_f \to 1} p_{\ell_f}(1-h_m, 1-h_f) = \infty$, and similarly for p_{ℓ_m}.[11]

Denote the utility function of an individual by u, where $u(c_i, p)$ is the utility from consuming private good c_i and public good p. We assume that u is twice continuously differentiable with $u_c > 0, u_p > 0, u_{cc} \leq 0, u_{pp} \leq 0$. We further restrict the class of utility functions below.

Both matching markets, the labor and marriage market, are competitive (full information and no search frictions) and there is no risk. The two markets and sorting choices therein are linked through the labor supply choice, which can be interpreted as a pre-labor market and post-marriage market continuous investment in 'effective' skills. This link is the key element of our model.

4.2 Decisions

The model is static and agents make three decisions, see Figure 3. In the *marriage market stage*, men and women choose their partner to maximize their value of being married. The outcome is a marriage market matching function, matching each woman x_f to some man x_m (or single-hood), and a market clearing price. In the second stage, the *household decision problem*, each matched couple chooses private consumption and allocates hours to the various activities—work in the labor market and at home—under anticipation of the labor market outcomes (matching function and wage function). This stage renders both private consumption and public consumption (and thus the time allocation), pinning down individuals' effective types. In the third stage, the *labor market stage*, agents take marriage market and household choices as given and match with firms based on their effective skills so that their wage income is maximized (or equivalently, each firm chooses an effective worker type to maximize profits). This problem pins down a labor market matching function and a market-clearing wage function.

In terms of exposition, we will describe the decision stages in reverse order.

[10] Our model can handle more general home production functions where part of the public good is purchased using wages. But given that we do not observe purchased public goods in the data below and given that spouses' observed home production time can be seen as producing a public good *net* of the purchased part, we focus on p as a function of hours only.

[11] Throughout, we will denote the partial derivatives of some generic function $f(x, y)$ using subscripts, for instance the derivative of f with respect to x is f_x unless there is risk of confusion, in which case we use $\partial f / \partial x$. We will denote the derivative of a function of a single argument by *prime*; and the derivative of any composition of functions using *brackets*, for instance, the derivative of $f(x, y(x))$ is denoted by $(f)_x$.

Labor Market. Taking marriage and household choices as given (in particular, the associated hours choices, which give rise to the *endogenous* distribution of effective types from which firms draw their workers, $\tilde{N}$), firms choose the *effective* worker type that maximizes their profits:

$$\max_{\tilde{x}} z(\tilde{x}, y) - w(\tilde{x}) \tag{1}$$

where $w : \tilde{\mathcal{X}} \to \mathbb{R}_+$ is the endogenous wage function taken as given. Market clearing pins down the labor market matching function $\mu : \tilde{\mathcal{X}} \to \mathcal{Y}$, mapping workers' effective skills to firm types in a measure-preserving way. Importantly, the labor market matching function μ depends on the hours choice (through $\tilde{N}$), which in turn will depend on the marriage partner. Thus, sorting on the two markets is connected.

And if $\tilde{\mathcal{X}}$ is an interval (as it will be the case below) then the first-order condition, which gives a differential equation in w, pins down the wage function as

$$w(\tilde{x}) = w_0 + \int_0^{\tilde{x}} z_{\tilde{x}}(t, \mu(t))dt, \tag{2}$$

where w_0 is the constant of integration. Note that $w(\tilde{x})$ is the wage of worker $\tilde{x}$ *per unit of time* (recall our time endowment is normalized to one unit and we will give a specific interpretation of a time unit when going to the data below), not the worker's earnings.

Household Problem. Each couple (x_f, x_m), taking the partner choice from the marriage market stage as given and anticipating the wage function and labor market matching function (w, μ), solves the following *cooperative household problem*. One partner (here w.l.o.g. the male partner) maximizes his utility subject to the household budget constraint and a constraint that ensures a certain level of utility for the female partner, by choosing the couple's private consumption and the hours allocation:

$$\begin{aligned} \max_{c_m, c_f, h_m, h_f} \quad & u(c_m, p(1-h_m, 1-h_f)) \\ s.t. \qquad & c_m + c_f - w(\tilde{x}_m) - w(\tilde{x}_f) = 0 \\ & u(c_f, p(1-h_m, 1-h_f)) \geq \overline{v}, \end{aligned} \tag{3}$$

where at this stage $\overline{v}$ is taken as a parameter by each household (it will be a function of female skills and endogenously determined in the next stage, the marriage market stage). When solved for all feasible $\overline{v} \in [0, \overline{v}_{max}(x_f, x_m)]$ (where $\overline{v}_{max}(x_f, x_m)$ is the maximum that x_f can obtain when matched with x_m), problem (3) traces out the household's pareto utility frontier. The solution to this problem yields the hours functions $h_i : \mathcal{X}_m \times \mathcal{X}_f \times [0, \overline{v}_{max}(x_f, x_m)] \to [0, 1]$ and consumption functions $c_i : \mathcal{X}_m \times \mathcal{X}_f \times [0, \overline{v}_{max}(x_f, x_m)] \to \mathbb{R}_+$. That is, for each partner in any matched couple (x_m, x_f) and for a given utility split $\overline{v}$, the problem pins down private consumption $c_i(x_m, x_f, \overline{v})$ and labor hours $h_i(x_m, x_f, \overline{v})$ (and therefore the public good $p(1-h_m, 1-h_f)$). Because the household problem is set up in a cooperative way, these allocations are pareto-efficient for any given wage function.

Marriage Market. Anticipating for each *potential* couple the solution to the household problem (h_f, h_m, c_f, c_m) as well as the labor market outcomes (μ, w), the value of marriage of man x_m from marrying woman x_f is given by the value of household problem (3) and thus:

$$\Phi(x_m, x_f, v(x_f)) := u(c_m(x_m, x_f, v(x_f)), p(1 - h_m(x_m, x_f, v(x_f)), 1 - h_f(x_m, x_f, v(x_f)))),$$

where we now make explicit that v, the marriage market clearing price, is an endogenous function of x_f and pinned down in the equilibrium of the marriage market. The marriage market problem for any man x_m is then to choose the optimal female partner type x_f by maximizing this value:

$$\max_{x_f} \Phi(x_m, x_f, v(x_f)). \tag{4}$$

The FOC of this problem (which gives a differential equation in v) together with marriage market clearing determine the marriage market matching function $\eta : \mathcal{X}_f \to \mathcal{X}_m$, mapping female skills to male skills in a measure-preserving way, and a transfer function $v : \mathcal{X}_f \to \mathbb{R}_+$, where $v(x_f)$ is the marriage payoff of woman x_f. The marriage matching function depends on the complementarities between men's and women's skills in Φ, as detailed below. Note that in principle, individuals can decide to remain single, which—given that there is an equal mass of men and women—will not happen in our baseline model if the value of marriage Φ is positive for all potential couples.

4.3 Equilibrium

We now formally define equilibrium.

Definition 1 (Equilibrium). *An equilibrium is given by a tuple $(\eta, v, h_m, h_f, c_f, c_m, \tilde{N}, \mu, w)$ such that*

1. *given (η, v, h_m, h_f), the pair (μ, w) is a competitive equilibrium of the labor market;*
2. *given (η, v, μ, w), the tuple (h_f, h_m, c_f, c_m) solves the household problem, pinning down $\tilde{N}$;*
3. *given $(\mu, w, h_m, h_f, c_f, c_m)$, the pair (η, v) is a competitive equilibrium of the marriage market.*

We next define a *monotone* equilibrium, which will be our main benchmark below.

Definition 2 (Monotone Equilibrium). *An equilibrium is monotone if it satisfies Definition 1 and:*

1. *labor market matching μ satisfies PAM, $\mu(\tilde{x}) = G^{-1}(\tilde{N}(\tilde{x}))$;*
2. *labor hours h_i are increasing in own type x_i and in partner's type x_j, $i, j \in \{f, m\}, i \neq j$, as well as in transfer v;*
3. *marriage market matching η satisfies PAM, $\eta(x_f) = N_m^{-1}(N_f(x_f))$, and v is increasing in x_f.*

In a monotone equilibrium, there are three additional requirements relating to the three different stages of this model. Specifically, 1. matching on the labor market is PAM; 2. hours worked in the labor market are increasing in own and in the partner's type; and 3. matching on the marriage market is PAM and the transfer to the wife is increasing in her type. Under 2. and 3., we obtain that a woman's effective type as a function of x_f, $\gamma_f(x_f) := e(x_f, h_f(\eta(x_f), x_f, v(x_f)))$, is strictly increasing in x_f since then $(\gamma_f)_{x_f} = e_{x_f} + e_h \left[\frac{\partial h_f}{\partial x_m}\eta' + \frac{\partial h_f}{\partial x_f} + \frac{\partial h_f}{\partial v}v' \right] > 0$ (where $(\gamma_f)_{x_f}$ denotes the total derivative of γ_f, see footnote 11, and where we denote by $h := h_i(\eta(x_f), x_f, v(x_f))$ the second argument of e), implying that γ_f can be inverted; and similarly for γ_m. As a result, in a monotone equilibrium, the endogenous cdf of effective types has a closed form, where the probability that $\tilde{x} \leq X$ is:[12]

$$\tilde{N}(X) = \frac{1}{2}N_f(\gamma_f^{-1}(X)) + \frac{1}{2}N_m(\gamma_m^{-1}(X)). \tag{5}$$

This discussion highlights an important point: The equilibrium hours function, h_f, not only depends on her own skill type x_f but also on marriage market outcomes: the skill type of her partner, $\eta(x_f)$, as well as the transfer guaranteed to her in the marriage $v(x_f)$; and similarly for men's hours function h_m. As a result, labor supply choices form the link between the marriage market (they are determined by the household and depend on who is matched to whom on the marriage market η) and the labor market (they affect the effective skill cdf $\tilde{N}$ and thus labor market matching μ and wages w).

This interdependence of marriage and labor market sorting is the crucial feature of our model. But it also makes the problem challenging from a theoretical point of view since we seek the simultaneous equilibrium of two intertwined matching markets, which are related through the time allocation choice. The analysis is further complicated by the possible feature of *imperfectly transferable utility* (ITU), where the hours functions and thus the public good production depend on transfer v.

To gain tractability and intuition into the main mechanisms of the model, we focus on a certain class of models (the quasi-linear class) that yields the *transferable utility* (TU) property, as explained below.

5 Analysis

In this section we show how the primitives of our model, in particular, of home production p and labor market production z, shape equilibrium.

5.1 The Quasi-Linear Class

We focus on a tractable specification of the household problem in which our model is TU-representable (see also Mazzocco, 2007 and Browning et al., 2014). The TU representation of our model obtains if the utility function falls into a known class, the *Gorman form*, which guarantees that utility is linear in the private good (possibly after a monotone transformation).[13] In our baseline model, we assume the

[12] To see this, observe that

$$\tilde{N}(X) = \frac{1}{2}\mathbb{P}[\gamma_f(x_f) \leq X] + \frac{1}{2}\mathbb{P}[\gamma_m(x_m) \leq X] = \frac{1}{2}\mathbb{P}[x_f \leq \gamma_f^{-1}(X)] + \frac{1}{2}\mathbb{P}[x_m \leq \gamma_m^{-1}(X)] = \frac{1}{2}N_f(\gamma_f^{-1}(X)) + \frac{1}{2}N_m(\gamma_m^{-1}(X)).$$

[13] More generally, the Gorman form of i's utility is given by $u^i(p, c_1, ..., c_n) = z^i(c_2, ..., c_n) + k(p)c_1$, which is *linear in at least one private consumption good*, with common coefficient $k(p)$, meaning that the marginal utility w.r.t. c_1 is equalized across partners, so that utility can be transferred between them at a constant rate.

utility function is quasi-linear in c:[14]

$$u(c_i, p) = c_i + p. \tag{6}$$

Then, the household's aggregate demand for private consumption c and public consumption p can be determined independently of the couple's sharing rule, v. As a consequence, the hours functions (h_f, h_m) are independent of v. In the marriage stage, in turn, the marital surplus is independent of the sharing rule v. As a result, the matching problem can be solved by maximizing the total value of marriage, independently of how it is shared (as in Shapley and Shubik, 1971 and Becker, 1973).

5.2 Conditions for Monotone Equilibrium

Our objective is to derive conditions under which any stable equilibrium is *monotone* in the sense of Definition 2. The monotone equilibrium will be our benchmark. We first state our main result and will then unpack its components to provide some technical details.

Proposition 1 (Monotone Equilibrium)**.** *If p is strictly supermodular and z is weakly convex in effective types $\tilde{x}$ and supermodular, then any stable equilibrium is monotone.*

The proof is in Appendix A.1.3. We call an equilibrium stable if it is robust to small perturbations (see Appendix A.1.1 for a precise definition). We show in that Appendix that if there exists an equilibrium there is at least one stable one.

A crucial condition for the monotone equilibrium is the home production complementarity (supermodular p, $p_{\ell_m \ell_f} > 0$). It gives rise to a 'progressive' way of organizing the household with gender balance in hours as opposed to specialization. If one partner works a lot in the labor market, then the other does as well, at the cost of less home production. The positive correlation of partners' hours within the household is clearly a force towards PAM in the marriage market: Having a partner with similar skills makes it easier to work similar hours in the labor market and, as the flip side, to put similar hours into home production, reaping the full benefits from the home production complementarity. Skilled individuals then prefer to match with skilled partners. In turn, positive sorting in the labor market stems from the complementarity between individuals' effective skills and jobs' skill requirements (supermodular z).

For the interested reader, we now unpack the technical details underlying Proposition 1, going over the three requirements of monotone equilibrium and why they are satisfied under the stated premise.

Positive Sorting in the Labor Market. As is well-known, if technology z is supermodular, then the worker-firm assignment in the labor market will satisfy positive sorting, that is the market-clearing matching function μ is increasing, where $\mu(\tilde{x}) = G^{-1}(\tilde{N}(\tilde{x}))$ is the firm matched to worker $\tilde{x}$.

Hours are increasing on own and partner's type. With quasi-linear utility (6), the household problem (3) takes the form:

[14] We here choose the simplest functional form in the Gorman class to reduce notation that obscures the main mechanism. In Online Appendix OB.1 we provide examples of other utility functions in this class under which our model features TU.

$$\max_{h_m,h_f} w(\tilde{x}_m) + w(\tilde{x}_f) - \overline{v} + 2p(1-h_m, 1-h_f), \tag{7}$$

where we substituted both the household's budget constraint and the wife's constraint to receive at least utility $\overline{v}$ into the objective function. As a consequence of TU, the overall split between public and private consumption (and thus the time allocation choice) can be made independently of how utility is shared, captured by $\overline{v}$. As a consequence, the hours functions (h_f, h_m) (and thus the public good) will only depend on types (x_f, x_m) but no longer on $\overline{v}$.

The FOCs of household problem (7) with respect to h_f and h_m are given by

$$w'(\tilde{x}_f)e_h(x_f, h_f) - 2p_{\ell_f}(1-h_m, 1-h_f) = 0 \tag{8}$$

$$w'(\tilde{x}_m)e_h(x_m, h_m) - 2p_{\ell_m}(1-h_m, 1-h_f) = 0. \tag{9}$$

In any interior solution for the hours choices of partners, each of these FOCs equalizes the marginal benefit of an additional hour in the labor market, captured by the wage gain, with its marginal cost stemming from a reduction in home production that affects both partners (hence the multiplication by 2).

To characterize under which conditions hours are increasing in own and partner's type, note that FOCs (8) and (9) give rise to two 'best-response' functions, one of wife's to husbands labor hours and one of husband's to wife's hours, whose intersection marks the equilibrium in the household decision stage. We are interested in how changes in (x_f, x_m) affect these hours choices, taking the equilibrium marriage market matching function η as given. In Appendix A.1.3, we show:

$$\frac{\partial h_f}{\partial x_f} = \frac{-(u)_{h_m h_m}(w)_{x_f h} + p_{\ell_m \ell_f}(w)_{x_m h}\eta'}{|H|} \tag{10}$$

$$\frac{\partial h_m}{\partial x_f} = \frac{-(u)_{h_f h_f}(w)_{x_m h}\eta' + 2p_{\ell_m \ell_f}(w)_{x_f h}}{|H|} \tag{11}$$

$$\frac{\partial h_m}{\partial x_m} = \frac{-(u)_{h_f h_f}(w)_{x_m h} + p_{\ell_m \ell_f}(w)_{x_f h}(\eta^{-1})'}{|H|} \tag{12}$$

$$\frac{\partial h_f}{\partial x_m} = \frac{-(u)_{h_m h_m}(w)_{x_f h}(\eta^{-1})' + 2p_{\ell_m \ell_f}(w)_{x_m h}}{|H|}. \tag{13}$$

The denominator in these expressions is given by $|H| = (u)_{h_f h_f}(u)_{h_m h_m} - 4p_{\ell_m \ell_f}$ (where the notation indicates this is the determinant of the Hessian of the household problem). By our definition of stability (Definition 3 in the Appendix A.1.1), $|H|$ is positive and moreover $(u)_{h_f h_f} \leq 0$ and $(u)_{h_m h_m} \leq 0$, derived from $u(w(\tilde{x}_f) + w(\tilde{x}_m) - \overline{v} + p, p) = w(\tilde{x}_m) + w(\tilde{x}_f) - \overline{v} + 2p(1-h_m, 1-h_f)$, see (7). Thus, in any stable equilibrium, (10)-(13) are strictly positive so that hours are increasing in own and partner's skill type if: home hours are complementary (p is supermodular), wages are supermodular in skills and hours, $(w)_{x_i h} = w'' e_{x_i} e_h + w' e_{x_i h} > 0$ where again $h = h_i(\eta(x_f), x_f)$ (guaranteed by convex and supermodular z), and if marriage market matching is PAM, $\eta' > 0$. We will specify primitives for which $\eta' > 0$ next.[15]

[15] Note that if (10)-(13) are strictly positive, then the distribution of effective types, $\tilde{N}$, can be pinned down in closed

Positive Sorting in the Marriage Market. Given the equilibrium hours functions (h_f, h_m), we obtain the value of marriage Φ as the value of the household problem (7):

$$\Phi(x_m, x_f, v(x_f)) = w(e(x_m, h_m(x_m, x_f))) + w(e(x_f, h_f(x_m, x_f))) - v(x_f) + 2p(1 - h_m(x_m, x_f), 1 - h_f(x_m, x_f)). \quad (14)$$

Complementarities among partners' types in Φ determine marriage market matching patterns. Under TU, $(\Phi)_{x_m x_f} = \Phi_{x_m x_f}$. If $\Phi_{x_m x_f} > 0$, then marriage matching is PAM, $\eta' > 0$.[16]

For consistency, we again adopt the male partner's perspective. Maximizing (14) with respect to x_f, while taking into account that $h_m(x_m, x_f)$ and $h_f(x_m, x_f)$ are already optimized so that they do not respond to further changes in x_f (by the Envelope Theorem), yields:

$$(\Phi)_{x_f} = 0 \qquad \Leftrightarrow \qquad w'(\tilde{x}_f) e_{x_f}(x_f, h_f(x_m, x_f)) - v'(x_f) = 0 \quad (15)$$

The transfer to the female partner, v, reflects the marginal impact of her type on her wage: When a man chooses a woman, he trades off the marginal benefits of choosing a higher type (which equals the marginal impact on her wage, $w' e_{x_f}$) with the marginal costs (which equals the marginal increase in transfer to her, v'). The higher the marginal wage return from a more productive female type, the larger is the increase in her compensation within the marriage. The reason why transfer v does not depend on the woman type's contribution to the public good p is that types only indirectly affect the public good production through the hours choice (and since hours were already optimized, the change in x_f has no impact on home production by the Envelope Theorem).

Then, the cross-partial derivative of Φ can be computed from (15) as:

$$\Phi_{x_m x_f} = (w)_{x_m x_f} = (w)_{x_f h} \frac{\partial \mathcal{h}_f}{\partial x_m} \quad (16)$$

highlighting that complementarities in the value of marriage must stem from the individual's wage depending on the partner's type through hours. $\Phi_{x_m x_f} > 0$ if wages are complementary in partners' skill types or, zooming in, when wages are supermodular in type and hours (meaning the marginal wage return to skill increases when putting in more labor hours) and when female labor hours are increasing in her partner's type, $\partial \mathcal{h}_f / \partial x_m > 0$, where $\mathcal{h}_f := h_f(x_m, x_f)$. Note that here, the comparative static of female hours with respect to male type is computed for *any* potential couple (x_m, x_f), not just for the ones that form in equilibrium $(\eta(x_f), x_f)$ (we still need to determine η at this stage), and we use the *mathcal*-notation in (16) to make this distinction from equation (13) clear. These sorting conditions are intuitive: There is PAM in the marriage market, so that x_f is matched to $\eta(x_f) = N_m^{-1}(N_f(x_f))$, if labor hours of spouses are complementary in the sense that an individual's labor hours are increasing

form and is given by (5). Moreover, $\tilde{\mathcal{X}}$ is a compact set in this case.

[16]Under ITU, $(\Phi)_{x_m x_f} > 0$ is equivalent to $\Phi_{x_m x_f} > \frac{\Phi_{x_f}}{\Phi_v} \Phi_{x_m v}$ (Legros and Newman, 2007). To see this, the FOC of problem (4) is given by $\Phi_{x_f} + \Phi_v v' = 0$, while $(\Phi)_{x_m x_f} = \Phi_{x_m x_f} + \Phi_{x_m v} v'$. Plugging the FOC into the latter condition gives the well-known Legros-Newman condition for PAM. In the quasi-linear class, this becomes $\Phi_{x_m x_f} > 0$ since $\Phi_{x_m v} = 0$.

in partner's type, and if at the same time working more hours boosts the marginal wage return to skill.

Using the formula for $\partial h_f/\partial x_m > 0$ (see Appendix A.1.3), we can re-express $\Phi_{x_m x_f}$ in a more symmetric way, which highlights the importance of the home production function also at the marriage stage:

$$\Phi_{x_m x_f} = 2p_{\ell_m \ell_f} \frac{(w)_{x_f h}(w)_{x_m h}}{|H|}. \tag{17}$$

Complementarity in home production (supermodular p) along with wages that are complementary in skill and hours (guaranteed by a convex and supermodular z) induce $\Phi_{x_m x_f} > 0$ and thus PAM in the marriage market.

5.3 Properties of Monotone Equilibrium and Stylized Facts

We now connect the properties of monotone equilibrium with our stylized facts in a *qualitative* way, before accurately replicating our facts in our quantitative analysis below.

Marriage Market Sorting. The property of positive sorting in the marriage market resembles our empirical finding of positive sorting on partners' education in Table 1.

Labor Market Sorting. In the monotone equilibrium, more skilled individuals work more in the labor market than at home compared to the less skilled, a feature that is reinforced by having more skilled partners. As a result, more skilled individuals have higher effective types, thereby obtaining more productive labor market matches: There is positive sorting in the labor market in (x, y), capturing the positive correlation between education and job's skill requirements in the data (Figure 1, left).

Marriage Market and Labor Market Sorting. The unique feature of our model is the link between labor and marriage market equilibrium and, in particular, labor and marriage sorting. This link becomes most transparent when highlighting how the labor market matching function depends on the marriage market matching function. Consider the total derivative $(\mu)_{x_f}$ (for $i = m$, this is similar), which—when positive—indicates PAM on the labor market in skills and skill requirements (x, y):

$$(\mu)_{x_f} = \mu' \left(e_{x_f} + e_h \left(\frac{\partial h_f}{\partial x_m} \eta' + \frac{\partial h_f}{\partial x_f} \right) \right). \tag{18}$$

Equation (18) illustrates how labor market sorting $(\mu)_x$ depends on marriage market sorting η'. When marriage market sorting is positive, $\eta' > 0$, then higher x are matched to higher y, $(\mu)_x > 0$ (given that hours of spouses are complementary $\partial h_f/\partial x_m > 0$). The intuition is straightforward. PAM on the marriage market induces individuals with higher x_i to have a better partner $x_j = \eta(x_i)$ and therefore to work more hours, which translates into a higher effective type $\tilde{x}_i$ and thus a better labor market match $y = \mu(\tilde{x}_i)$, compared to when marriage market sorting is not positive. In a stylized way, this property of the monotone equilibrium is related to our empirical fact that labor market sorting is stronger for positively sorted couples (Figure 1, right).

The Role of Hours. In the monotone equilibrium, labor hours are complementary within couples: Increasing, say, female skills, not only pushes up her own labor hours but also induces her partner to work more. As a result, partners' hours co-move. There are two drivers behind this result. First, for a *given male partner type* x_m (for exogenous marriage matching), an increase in female skills increases her labor hours. But this reduces her home hours, inducing her partner to also work less at home and more in the market due to $p_{\ell_m \ell_f} > 0$. As a result, *both* partners increase their labor hours as the female skill improves. Second, this complementarity is reinforced under *endogenous* marriage market sorting: Under PAM, an increase in her skill x_f leads to a better partner $x_m = \eta(x_f)$, who by himself puts in more labor hours and less home hours. And since $p_{\ell_m \ell_f} > 0$, the wife adjusts hours in the same direction (less home hours and more labor hours), reinforcing the co-movement of hours within the couple. Thus, PAM on the marriage market fuels the complementarity of hours within couples—a feature we saw in the data (Figure 2).

Finally, an interesting feature of the monotone equilibrium is that it can be consistent with a gender gap in labor market sorting: If the home production function is such that women spend relatively more time at home (e.g. if they are relatively more productive at home), then men will be 'better' matched on the labor market compared to women of the same skill. Thus, our competitive model can generate a gender gap in sorting and wages even in the absence of discrimination or differential frictions.

To see this, consider labor market sorting in terms of firm productivity and skills (x_i, y), and how it varies across gender $i \in \{f, m\}$. Consider a man and a woman with $x_f = x_m$. We say that x_m is 'better sorted' than x_f if $\mu(e(x_m, h_m)) > \mu(e(x_f, h_f))$. For each man and woman of equal skills, $x_f = x_m$, men x_m is better sorted if he works more hours on the labor market, $h_m(x_m, \eta^{-1}(x_m)) > h_f(\eta(x_f), x_f)$, which will help rationalizing our finding in the data on the differential sorting of men and women in the labor market (solid lines Figure 1, left). But controlling for hours worked, $h_m(x_m, \eta^{-1}(x_m)) = h_f(\eta(x_f), x_f)$, closes the sorting gap in the model and considerably shrinks it in the data (dashed lines Figure 1).

5.4 Non-Monotone Equilibrium

The monotone equilibrium captures—albeit in a stylized way—several salient features of the data. Some features of the monotone equilibrium, in particular the complementarity of spouses' hours, may be in contrast to the traditional and more standard view of the household, which relies on specialization. Historically, it is plausible that a different equilibrium was in place, in which partners' hours in home production were substitutable and where positive sorting on the marriage market was less pronounced or sorting was even negative, giving rise to specialization of household members. We capture this different regime by an equilibrium that—with some abuse—we call *non-monotone equilibrium* and we highlight the role played by properties of the home production function. We define a non-monotone equilibrium as the monotone one with two differences. First, there is negative assortative matching (NAM) in the marriage market. And second, labor hours are *decreasing* in partner's type.

Proposition 2 (Non-Monotone Equilibrium)**.** *If p is strictly submodular and z weakly convex in effective types $\tilde{x}$ and supermodular, then any stable equilibrium is non-monotone.*

The proof is in Appendix A.1.4. This result highlights the key role of home production complementarities/substitutabilities in shaping equilibrium. Making hours at home substitutable, $p_{\ell_m \ell_f} < 0$, gives rise to an equilibrium that relies on 'specialization', where a more skilled partner puts more labor hours while own labor hours go down in response. At the same time, the partner spends less time in home production while own home production time increases. This specialization within the household is clearly a force towards NAM in the marriage market, which indeed materializes. The reason is that increasing the partner's type pushes own labor hours *down*, hurting own labor market prospects especially for skilled individuals. Skilled individuals then prefer to match with less skilled partners.

The only feature that both equilibria have in common is PAM on the labor market not only in $(y, \tilde{x})$ (guaranteed by supermodular z) but, importantly, also in (y, x), which follows from equation (18).[17]

Thus, complementarity vs. substitutability of home hours shapes equilibrium. In particular, $p_{\ell_m \ell_f} \lesseqgtr 0$ determines whether marriage partners match positively and whether their hours—both at home and at work—are complementary. The monotone equilibrium captures 'progressive' times while the non-monotone one reflects a 'traditional' division of labor. To our knowledge, this mechanism in which home production complementarities are the key determinant of both marriage and labor market outcomes is new in the literature. We now investigate the nature of home production and our mechanism in the data.

6 Quantitative Model

One advantage of our parsimonious model is that we obtain analytical properties that illuminate its mechanism. To evaluate its quantitative importance, we now augment our model so that it can match the data. We do so by implementing minimal departures from our baseline model to preserve its mechanism.

6.1 Set-Up and Decisions

Our objective is to build a quantitative version of our baseline model that can match key facts of the data while minimally departing from our original set-up. To this end, we augment the model by including shocks in each of the three stages—marriage market, household decision stage, and labor market—so that we capture the following: imperfect sorting and non-participation on both marriage and labor markets as well as heterogeneity in hours choices across couples of the same type. Importantly, we show in Proposition O1 (Online Appendix OB.2) that under similar conditions as in the baseline model, the properties of monotone equilibrium hold *on average* in our augmented model. We make three changes:

First, in order to capture mismatch in the labor market along (x, y), we augment individuals' education/skill x by an idiosyncratic productivity component ν. We assume that individuals are characterized by *discrete* human capital $s := k(x, \nu) \in \mathcal{S}$, distributed according to cdf N_s, where s takes the role of x from the baseline model. We assume ν (and thus s) is observed by the agents in the market, but not by us. In the labor market, the match relevant attribute of a worker is her effective human capital $\tilde{s} := e(s, h)$ (instead of $\tilde{x}$), whose distribution we denote by $\tilde{N}_s$. The firm now solves: $\max_{\tilde{s}} z(\tilde{s}, y) - w(\tilde{s})$.

[17] Further note that the distribution of effective types $\tilde{N}(X) = \frac{1}{2}N_f(\gamma_f^{-1}(X)) + \frac{1}{2}N_m(\gamma_m^{-1}(X))$ will be pinned down just like in the monotone equilibrium since own labor hours are still increasing in own type so that γ_i is still invertible.

Second, we account for heterogeneity in labor supply *within* (s_f, s_m)-type couples and *within* s_i-type singles (and for non-participation) by introducing idiosyncratic labor supply shocks. We denote by δ^{h_i} the idiosyncratic preference of an agent for hours alternative $h_i, i \in \{f, m\}$. In this quantitative version of our model, hours are discrete elements of choice set $\mathcal{H}$, $h_i \in \mathcal{H}$. Each decision-maker (single or couple) draws a vector of labor supply shocks, one for each alternative h_i. These shocks realize after marriage.

In the household decision stage, partners now maximize utility *plus* labor supply shock:

$$\begin{aligned} \max_{c_m, c_f, h_m, h_f} \; & u(c_m, p^M(1-h_m, 1-h_f)) + \delta^{h_m} \\ s.t. \quad & c_m + c_f - w(\tilde{s}_m) - w(\tilde{s}_f) = 0 \\ & u(c_f, p^M) + \delta^{h_f} \geq \overline{v}. \end{aligned} \tag{19}$$

where we introduce the notation p^M for the home production technology of couples (Married).

Similarly, the consumption-time allocation problem of singles is given by

$$\begin{aligned} \max_{c_i, h_i} \quad & u(c_i, p^U(1-h_i)) + \delta^{h_i} \\ s.t. \quad & c_i - w(\tilde{s}_i) = 0 \end{aligned} \tag{20}$$

where we denote by p^U the home production function of singles (Unmarried).

Third, to accommodate the fact that marriage market matching on *human capital* s may not be perfectly assortative and to account for non-participation/single-hood, we introduce an idiosyncratic taste shock for partners' s-types. We denote by β_m^s and β_f^s the idiosyncratic taste of man m and woman f for a partner with human capital $s \in \{\mathcal{S} \cup \emptyset\}$ where $s = \emptyset$ indicates the choice of remaining single. Each individual draws a vector of taste shocks, one for each discrete alternative s. So, individuals in the marriage market value potential partners not only for their impact on the economic joint surplus (as before) but also for their impact on the non-economic surplus (which depends on preference shocks β_f^s or β_m^s). The marriage problem of a man with s_m now reads

$$\max_s \quad \Phi(s, s_m, v(s)) + \beta_m^s$$

where the choice of marrying a woman of any human capital type $s = s_f$ needs to be weighed against the choice of remaining single $s = \emptyset$ ($\Phi(\emptyset, s_m, v(\emptyset))$ is the economic value of remaining single).

Similar to the baseline model, Φ captures the *economic* surplus from marriage. Different from the baseline model, due to the introduction of labor supply shocks that have not yet realized at the time of marriage, Φ is the *expected* economic surplus from marriage. The expectation is taken over the different hours alternatives of the couple whose choice probabilities are pinned down at the household stage. See Appendix B for details. Since marriage market matching is no longer pure (due to both the discreteness of the match attribute s and the idiosyncratic shocks β^s), $\eta : \{\mathcal{S} \cup \emptyset\}^2 \rightarrow [0, 1]$ here denotes the matching *distribution* (as opposed to the matching function).

6.2 Functional Forms

We parameterize our model as follows. The production function on the labor market is given by $z(\tilde{s}, y) = A_z \tilde{s}^{\gamma_1} y^{\gamma_2} + K$ where A_z is a TFP term, (γ_1, γ_2) are the curvature parameters reflecting the elasticity of output with respect to skill and firm productivity, and K is a constant.

For couples, the public good production function is assumed to be CES

$$p^M(1-h_m, 1-h_f) = A_p\Big[\theta\big(1-h_f\big)^\rho + (1-\theta)\big(1-h_m\big)^\rho\Big]^{\frac{1}{\rho}}$$

where A_p is the TFP in home production, θ is the relative productivity of a woman, and ρ is the parameter that determines the elasticity of substitution, $\sigma := 1/(1-\rho)$ (where $\sigma < (>)1$ indicates that spouses' home hours are strategic complements (substitutes)). We assume that home production for singles is given by $p^U(1-h_i) = A_p\Theta_i\big(1-h_i\big)$ where $\Theta_i \in \{\theta, 1-\theta\}$ depending on gender.

The utility function of individual i is given by $u(c_i, p) = c_i + p$ where $p \in \{p^M, p^U\}$ for spouses and singles, and where we assume that both men and women have the same preferences. We adjust the private consumption of singles by the McClemens factor (Anyaegbu, 2010).

Human capital as a function of skill and productivity shock is given by $s \propto x + \nu$, where we assume that s is proportional to the sum of observed skill and (to us) unobserved productivity.

We specify the effective human capital functions as:

$$\begin{aligned} \tilde{s}_f &= \psi s_f h_f \\ \tilde{s}_m &= s_m h_m \end{aligned} \tag{21}$$

where, if a man and a woman have the same (s, h)-combination, $\tilde{s}_f \leq \tilde{s}_m$ if $\psi \leq 1$. We thus allows for a labor market penalty for women that could reflect gender discrimination or productivity differences.

Finally, both marriage taste shocks and labor supply shocks follow extreme-value type-I distributions:

$$\begin{aligned} \beta^s &\sim \text{Type I}(\bar{\beta}^t, \sigma^t_\beta) \quad && \text{for } t \in \{M, U\} \text{ and } s \in \ \{\mathcal{S} \cup \emptyset\} \\ \delta^{h^t} &\sim \text{Type I}(\bar{\delta}, \sigma_\delta) \quad && \text{for } t \in \{M, U\} \text{ and } h^t \in \ \mathcal{H} \end{aligned}$$

where we allow for different preference shock distributions for marriage partners and singles—index t indicating the household type.[18] We normalize the location parameter of both labor supply and marriage market preference shocks to zero, $\bar{\delta} = \bar{\beta}^M = \bar{\beta}^U = 0$. Moreover, we specify the labor supply shocks as

$$\delta^{h^t} = \begin{cases} \delta^{h_i}, i \in \{f, m\} & \text{if } t = U \\ \delta^{h_f} + \delta^{h_m} & \text{if } t = M. \end{cases}$$

[18] Note that without the different scales for partner and single choices our parsimonious model (featuring no couple/single-specific parameters) would have difficulty to generate enough singles. Allowing for different scales, however, means that our marriage market resembles a nested logit problem with degenerate (single) nest, associated with known identification issues for the scale of the degenerate nest (Hunt (2000)). This is why we fix σ^U_β outside of the main estimation below.

That is, when making hours choices, a decision-making unit draws *a single* labor supply shock, δ^{h^t}, that is extreme value distributed. In the case of singles, the decision-making unit is just one person and hence, as is standard, this agent draws a shock for each hours alternative. In the case of spouses however, the decision-making unit is the couple. Therefore, that household draws a single shock for each *joint* time allocation of the spouses (equivalently, the sum of the spouses' shocks is assumed to be extreme-value distributed). We make this adjustment to the standard setting, where *each* individual agent draws an extreme-value shock when making a discrete choice, in order to obtain tractable choice probabilities for the joint hours allocation that help with computation and identification of the model.[19]

6.3 Model Solution

Appendix B describes the numerical solution of the quantitative model in detail. It consists of solving a fixed point problem in the wage function w (or, equivalently, in the hours functions (h_f, h_m)). For any given wage function, agents make optimal marriage and household choices as well as labor market choices. Labor market choices then give rise to a new wage function that, in equilibrium, needs to coincide with the initially postulated wage function. We implement a search algorithm that iterates between the problem of households and firms, producing a new wage function at each round, and that halts when the wage function satisfies a strict convergence criterion. Our procedure ensures that at convergence, both the labor *and* the marriage market are in equilibrium and households act optimally.

A challenge in our fixed point algorithm is that when partners determine whether a particular hours choice is optimal (which—as discussed—can be understood as an 'investment' in effective skills), they must compare the payoff of this investment with all alternative investments. But the competitive wage only determines the price for *equilibrium* investments.[20] In order to obtain the off-equilibrium wages without significantly perturbing the equilibrium wages, we use a *tremble strategy*. We postulate that a small fraction of agents are tremblers who make a mistake by choosing off-equilibrium hours. This ensures that also off-equilibrium choices will be priced and individuals can compare *all* investment choices when solving the household problem. While trembling is a widely used concept in game theory, we believe the application to matching markets with investment is new.

7 Estimation

We estimate our model in order to assess whether partners' home production time is complementary or substitutable in the data; and how this home production property (and how it evolved over time) shapes empirical sorting patterns on both marriage and labor market, and ultimately household income inequality as well as gender disparities in labor market outcomes in the cross-section and over time.

[19] Gayle and Shephard (2019) follow a similar approach in their numerical solution (see their footnote 24), where they assume households draw one shock for each of the couple's joint hours-combinations.

[20] This issue is similar to the one in Cole, Mailath, and Postlewaite (2001) who study bargaining in a matching problem with pre-match investment. Apart from trembling, the hours shocks also help us to price all investment alternatives.

7.1 The Data

We again use data from the German SOEP combined with information from the dataset of occupational characteristics (BIBB). The challenge is to bridge our static model with the panel data which is intrinsically dynamic and contains life-cycle features. We deal with it as follows. For the estimation of worker unobserved heterogeneity (which will be done outside of the model), we exploit the full panel structure in order to make use of techniques that control for unobserved time-invariant characteristics. In turn, for the structural estimation of the model we construct a dataset that features each individual only once while accounting for his/her 'typical' outcomes. To be able to assess the typical outcomes, we focus in our baseline analysis on a restricted time period (2010-2016) so that each individual is captured in only one life-cycle stage and we focus on observations that are not too different in age (25-50). We consider each individual as one observation and generate summary measures (or 'typical' outcomes) of the life-cycle stage we see them in. We then define for each individual the typical occupation (based on a combination of tenure and job ladder features), typical labor hours and typical wage in that occupation, and typical home hours while holding that occupation, as well as the typical marital status. In line with our model, we only consider those individuals who are either married/cohabiting or have never been married and are thus single. We drop divorced and widowed people because they likely behave differently than the singles in our model. Our final sample contains 5,153 individuals, 50% of which are men. In Online Appendix OC.2, we provide the details of the sample construction.

7.2 Identification

We need to identify 10 parameters and two distributions. We group the parameters into 5 categories and discuss the identification group-wise. We have parameters pertaining to the home production function (θ, ρ, A_p), the production function $(\gamma_1, \gamma_2, A_z, K)$, labor supply and marriage preference shocks distributions $(\sigma_\delta, \sigma_\beta^M)$, and a labor productivity wedge (ψ). Finally, we have the distributions of worker human capital and job productivity (N_s, G). We provide formal identification arguments in Appendix C.1 and summarize the logic here. Our estimation will mostly be parametric. Nevertheless, we consider it useful to lay out non-/semi-parametric arguments in order to understand the source of data variation that pins down our parameter estimates. We will also clarify which parametric restrictions (mainly pertaining to the shock distributions) are important.

The home production function, and thus (θ, ρ, A_p), is identified from choice probabilities for home hours by households of different s-types. The formal identification uses the assumption that the labor supply shock for different hours choices of husband and wife follows a type-I extreme-value distribution.

The production function, and thus $(\gamma_1, \gamma_2, A_z, K)$, is identified from wage data. In our competitive environment, there is a tight link between wages and the marginal product (and thus technology), which allows us to do so. The curvature and TFP parameters, $(\gamma_1, \gamma_2, A_z)$, can be identified from the first derivative of the wage function (the marginal product), following arguments from the literature on the identification of hedonic models (Ekeland, Heckman, and Nesheim, 2004). In turn, the constant in the production function (K) can be identified from the minimum observed hourly wage.

The pair $(\sigma_\delta, \sigma_\beta^M)$ associated with our shock distributions is identified as follows (note that in each distribution we make one normalization choice). In the absence of labor supply shocks, any two couples of the same type (s_f, s_m) would choose the same combination of hours. Hence, the variation in hours choices *by couple type* pins down the scale parameter of the labor supply shock distribution σ_δ. Similarly, in the absence of any preference shocks for marriage partners ($\sigma_\beta^M = 0$), the model would produce perfect assortative matching on the marriage market with $corr(s_f, s_m) = 1$. The extent of marriage market sorting and mismatch identifies the scale parameter of preference shocks for partners, σ_β^M. Note that the standard result in the literature that the scale parameter is not identified separately from the utility associated with the discrete choices (e.g. Keane, Todd, and Wolpin, 2011) does not apply in our context. The reason is that we are able to identify utility in a prior step from household labor supply choices. Importantly, we do not exploit variation in partner choices to identify the utility and therefore, this variation can be used to identify the scale of the marriage shock distribution. Our identification result relies on the extreme-value assumption of the shock distributions, yielding tractable choice probabilities.

The productivity or discrimination wedge of women, ψ, is identified by the hourly gender wage gap *conditional on hours and s-type.* If there was no wedge, $\psi = 1$, women and men with the same (s, h)-bundle should receive the exact same wage. A gap can only be rationalized by $\psi \neq 1$.

Finally, the worker and job heterogeneity will be identified directly from the data. We use the empirical distributions of workers' human capital and occupations' productivity for (N_s, G). In sum:

Proposition 3 (Identification)**.** *Under the functional form assumptions from Section 6.2 and Assumption D1 (Appendix), the model's parameters are identified.*

Our identification result informs the moments we choose to pin down our parameters. To identify the home production function, we use five moments related to the division of labor and to the complementarity of hours within households (ratio of labor force participation of women to men; ratio of labor force participation of married to single individuals, by gender; ratio of full-time work of women to men, and correlation of spouses' home production hours). To identify the production function, we use four moments related to the hourly wage distribution (its mean, variance, and the 90-10 and 90-50 percentiles). To identify the marriage shock parameter we use two marriage market moments (the correlation of spouses' human capital types and fraction of single men). To identify the scale of the labor supply shock, we use four moments related to the hours variation across households of given human capital (female labor force participation rate by couple type and single type, where we select 2 types). Finally, we identify the female labor wedge with two moments related to the gender wage gap conditional on (s, h). In total we have 17 moments, described in detail in Table O4 in Online Appendix OD.

7.3 Two-Step Estimation

We propose a two-step estimation procedure. The first step estimates worker and job heterogeneity as well as the constant in the production function *outside* of the model. In a second step, given the worker and job distributions, we estimate the structural parameters of the primitives *within* the model.

7.3.1 First Step: Estimation Outside the Model

In a first step, we estimate worker types s (or (x, ν)) and job types y. Except for x (education), these types are not directly observed. Moreover, even though we observe the educational group of a worker we need to translate it into productivity units.

ESTIMATION OF WORKER TYPES. Let $ed \in \{hs, voc, c\}$ be the education level of an individual (standing for *h*igh *s*chool or less, *voc*ational training and *c*ollege) and ν be their ability. Based on our theory, we specify an empirical model for hourly wages, namely as a function of effective types (which in turn are a function of education, ability and hours worked). That is, we assume the empirical log hourly wage of individual i at time t (where t is a year in our sample) is given by

$$\ln w_{it} = \nu_i + \sum_{ed \in \{voc, c\}} \alpha^{ed} x_{it}^{ed} + \beta_1 h_{it} + \beta_2 h_{it}^2 + \beta_z' Z_{it} + \kappa_s + \rho_t + \epsilon_{it} \tag{22}$$

where x_{it}^{ed} are indicator variables for the education group of an individual (meant to capture x in our model) at time t. Coefficient α^{ed} gives the 'value' of education ed in terms of log wage units, where $0 < \alpha^{voc} < \alpha^c$ would indicate positive returns to education. While these coefficients indicate the average return to education for all individuals in a certain category, ν_i is a person fixed effect capturing unobserved time-invariant ability, with model counterpart ν. In turn, h_{it} denotes (typical) weekly labor hours (capturing the time 'investment' in labor productivity in our model). Finally, Z_{it} are time-varying controls for the individual, κ_s and ρ_t are state and time fixed effects, and ϵ_{it} is a mean-zero error term.[21]

We thus make use of the dynamic features of the (panel) data to estimate individual unobserved heterogeneity ν. For computational tractability, we divide individuals in each education bin into two groups depending on their ν_i (above and below the median). We compute the level of the low and the high ability based on the average fixed effect in their group, so $\nu_i \in \{\nu_L, \nu_H\}$. Hence, individuals belong to one of six human capital bins (three education types times two ability types). We then order individuals by their human capital $s_i = \alpha^{ed} x_i^{ed} + \nu_i$, giving us a global ranking of worker types. We use the empirical cdf of s_i as our estimate for workers' human capital distribution N_s.

There are three challenges in implementing (22): First, there may be confounding factors impacting both hours and wages. While we deal with time-invariant unobserved heterogeneity using the panel regression with individual fixed effects, time-varying unobserved heterogeneity, such as productivity

[21] We do not include occupation fixed-effects since in our model, conditional on $\tilde{s}$ (which we control for here by controlling for (x, ν, h)), the wage does not depend on occupation in our competitive equilibrium. But even doing so—which we have done for robustness—does not significantly change the impact of x or ν on the hourly wage.

shocks or health shocks, could still be problematic. To address this concern, we use an instrumental variable (IV) approach. In our model, there is a systematic relationship between the hours worked of an individual and the hours worked by their partner, so we use the partner's hours as an instrument for own hours. Identification relies on changes in spousal labor hours over time. The identifying assumption is that conditional on the individual fixed effect and education, partner's hours are exogenous in the wage regression and that partner's labor hours impact own wage *only* through own labor hours, which is satisfied in our model. Second, we only observe wages for those who work and labor market participation is not random. To account for selection, we apply a Heckman selection correction (Heckman, 1979). Third, even when we control for selection, using (22) we can only estimate types for those individuals who are employed for at least two periods in our panel. We therefore impute the fixed effects for those who we never observe participating using the *multiple imputation* method.

We provide the details on the sample as well as on the IV, selection and imputation in Appendix C.2. The estimation results of (22) are in Table 6 and the estimated skill distribution, in Table 7. Note that regression (22)—under the IV approach—delivers a causal effect of hours on hourly wages, where we find that increasing weekly hours worked from 30 to 40 increases hourly wages by around 4%.[22]

Estimation of Job Types. The empirical counterpart of our model's firms are occupations (we do not observe firms in the GSOEP). As in Section 3, we measure occupations' productivity types y from data on their task complexity. Our main dataset is the BIBB (comparable to the O*NET in the US), giving extensive information on task use in each occupation, were we focus on 16 tasks measured on a comparable scale. We measure the occupations' types in two steps. First, we use a Lasso wage regression to select the important/pay-off relevant tasks. In a second step, we run a principal component analysis (PCA) to reduce the task dimensions further to a single one, where we use the (normalized) first principal component as our one-dimensional occupation characteristic y. Importantly, we use the wage regression only to select the relevant tasks but we do not use the estimated coefficients. See Appendix C.3 for the details of this approach and for the alternative approaches that we pursued for robustness and which have led to similar results.

Estimation of Constant in Production Function. We assume that the constant in the production function is not shared between workers and firms but accrues to the worker in form of a minimum hourly wage (the wage of someone with the lowest human capital who will be matched to the lowest productive occupation, $y = 0$). This way, we obtain $K = 6.32$.

7.3.2 Second Step: Internal Estimation

There are nine remaining parameters of the model, $\Lambda \equiv (\theta, \rho, A_p, \gamma_1, \gamma_2, A_z, \psi, \sigma_\delta, \sigma_\beta^M)$. They are disciplined by 17 moments that we chose based on our identification arguments (Section 7.2). To estimate

[22] Our estimated effect on our sample of men and women in Germany is smaller but comparable to effects estimated on US data: Aaronson and French (2004) also use a panel regression with fixed effects with an IV for hours and finds that increasing hours from 20 to 40 per week increases the hourly wage by 25%; Bick et al. (2020) (who focus on men) find that increasing hours from 30 to 40 per week increases hourly wages by 11%.

these parameters, we apply the method of simulated moments (McFadden, 1989; Pakes and Pollard, 1989). For any vector of parameters, Λ, the model produces the 17 moments, $mom_{sim}(\Lambda)$, that will also be computed in the data, mom_{data}. We then use a global search algorithm to find the parameter values that minimize the distance between simulated and observed moments. Formally, the vector $\hat{\Lambda}$ solves

$$\hat{\Lambda} = \arg\min_{\Lambda} \quad [mom_{sim}(\Lambda) - mom_{data}]'\mathcal{V}[mom_{sim}(\Lambda) - mom_{data}]$$

where $\mathcal{V}$ is specified as the inverse of the diagonal of the covariance matrix of the data.

7.4 Results and Fit

We report the parameters that we fixed outside of the structural estimation, $(K, \bar{\delta}, \bar{\beta}^M, \bar{\beta}^U, \sigma^U)$, in Table 10 (Appendix C.4.1). The estimated parameters are in Table 2. The result we want to highlight is that our estimates indicate that spouses' hours at home (and therefore, in the labor market) are complements with $\rho = -0.54$, pushing the model towards the *monotone equilibrium* of the baseline model. The main data moment calling for a negative ρ is the strong positive correlation of spouses' home hours.[23] Further, the estimated home production function indicates that women are significantly more productive at home than men ($\theta = 0.78$). The large differences in labor force participation and full time work across gender call for this relatively high female productivity at home. In terms of labor market production, our estimates indicate that it is concave in both the workers' effective skill as well as the jobs' productivity ($\gamma_1 < 1, \gamma_2 < 1$). Labor market TFP A_z is estimated to be higher than home production TFP, A_p. The empirical gender wage gap *conditional* on hours and human capital calls for a female productivity/discrimination wedge, which we estimate as $\psi = 0.84$. This implies that, for any given type and choice of hours, women's effective skills are 16% lower than those of men.

Finally, regarding the marriage preference and labor supply shocks, our estimated scale parameters ensure that we match the fraction of singles, the extent of mismatch in the marriage market, and the heterogeneity in hours choices by households of the same human capital type. We also report the standard errors of the estimates.[24] The last column presents our sensitivity analysis (Andrews, Gentzkow, and Shapiro, 2017) where we report the most important moments that explain 50% of the impact on each parameter in estimation.[25] Our sensitivity analysis is in line with our identification arguments. For example, the correlation of spouses' hours, M_5, is an important moment disciplining the home production complementarities, ρ; or, the female productivity wedge, ψ, is most related to the within-type gender wage gaps, M_{12} and M_{13}.

[23] Further, if it was the case that $\rho > 0$ (i.e. if hours were strategic substitutes), then marriage market sorting would be random, see Figure 7d where marriage sorting drops significantly as ρ becomes positive.

[24] The covariance matrix of the estimator is computed as matrix $Var = \left[D_m' \mathcal{V} D_m\right]^{-1} D_m' \mathcal{V} C \mathcal{V}' D_m \left[D_m' \mathcal{V} D_m\right]^{-1}$, where D_m is the 10×17 matrix of the partial derivative of moment conditions with respect to each parameter at $\Lambda = \hat{\Lambda}$ and C is the covariance matrix of the data moments.

[25] We compute the sensitivity of each parameter to the moments as $|Sensitivity| = |-\left[D_m' \mathcal{V} D_m\right]^{-1} D_m' \mathcal{V}|$, defined by Andrews et al. (2017), see footnote 24 for notation.

Table 2: Estimated Parameters

Parameter	Estimate	s.e.	Top Sensitivity Moments
Female Relative Productivity in Home Production, θ	0.78	0.02	M11, M2, M9
Complementarity Parameter in Home Production, ρ	-0.54	0.22	M5, M3, M13
Home Production TFP, A_p	41.38	0.98	M11
Elasticity of Output w.r.t. $\tilde{s}$, γ_1	0.59	0.05	M2, M11, M8, M9
Elasticity of Output w.r.t. y, γ_2	0.16	0.07	M2, M8, M11
Production Function TFP, A_z	42.33	2.28	M2, M11, M8
Female Productivity Wedge, ψ	0.84	0.03	M13, M12
Labor Supply Shock (scale), σ_δ	7.51	0.40	M2, M11
Preference Shock for Partners (scale), σ_β^M	0.19	0.02	M11, M2, M13, M10

Notes: s.e. denotes Standard Errors. *Top Sensitivity Moments* reports the most important moments explaining 50% of the total impact on each parameter in estimation, based on our sensitivity measure (see footnote 25). $M1, ..., M17$ denote the 17 targeted moments (see Table 11, Appendix C.4).

Figure 4 summarizes the fit between model and data moments, where we plot all 17 moments (red dots indicate the level of these moments in the model) as well as their blue confidence interval of the corresponding data moment (computed from a bootstrap sample). We re-scaled some moments ($M_6 - M_9$) to be able to plot them all in the same graph. Table 11 in Appendix C.4 reports the fit in detail and indicates the moments corresponding to numbers 1-17. Our model achieves a good fit with the data, with nearly all model moments lying in the confidence interval of their data moments.

Figure 4: Model Fit: Model Moments (red) with Data Confidence Intervals (blue)

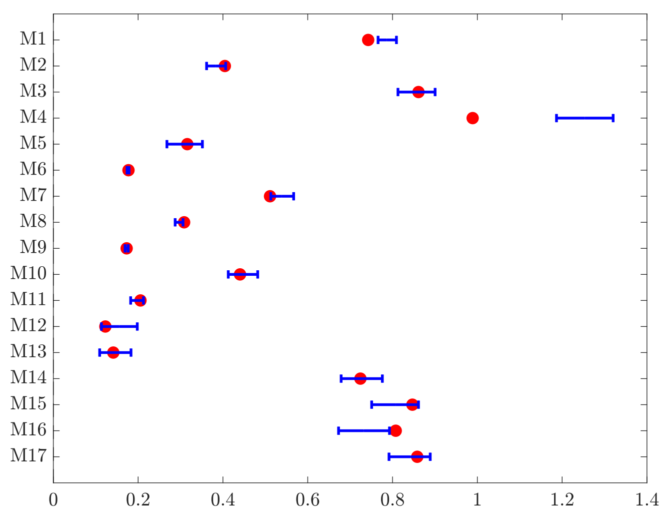

7.5 Model Validation

Apart from fitting the aggregate moments targeted in estimation, our model reproduces rich, un-targeted features: The relation between marriage and labor market sorting, and the link (hours) between them that we documented in Section 3.2.

Marriage Market Sorting. Table 12, Appendix C.4, displays the matching frequency of marriages by three education types (low, medium, and high) in data and model. The main panel indicates the frequencies of different types of couples while the bottom row (right column) indicates the frequencies of single men (women) by education. Data frequencies are in parentheses. In our estimation, we only targeted the *overall* correlation of couples' *human capital* types (i.e. s-types), as s is the relevant matching characteristic on the marriage market in our model. We did not target marital matching on *education*, x, especially not the detailed matching frequencies. Nevertheless, the model matches well the observed marriage frequencies by education type: A considerable fraction of couples matches along the diagonal, while the off-diagonal cells indicate that mixed couples (especially high-low couples) are rare—a sign of positive assortative matching on education. Our model also captures that medium educated men and women are most likely to be single, where we only targeted the average fraction of male singles.

Labor Market Market Sorting. We report in Figure 5, left panel, the labor market matching function for men (blue) and women (red) in the model (solid) and data (dashed). It is given by job productivity y as a function of individuals' human capital s. Our model captures that labor market sorting is PAM and that men are better matched for any given level of human capital.

Relationship between Labor Market Sorting and Marriage Market Sorting. We documented in Section 3.2 a strong link between labor market and marriage market sorting in the data, where labor market sorting is maximized for individuals who are well matched in the marriage market. Figure 6 (left panel), which compares data and model, shows that our model reproduces this pattern. Note that consistent with our quantitative model (and in contrast to Section 3.2), we here proxy marriage market sorting by spouses' differences in *human capital* s-types (as opposed to differences in education), also in the data. Similarly, labor market sorting is measured by the correlation of (s, y) (instead of (x, y)).

Hours as the Link between Marriage Market Sorting and Labor Market Sorting. The key feature of our model is that marriage and labor markets are linked in equilibrium, namely through the household's time allocation choice. Here we show that the model replicates salient features of the data according to which hours are associated with both marriage and labor markets outcomes. Figure 6, right panel, shows that both in data (dashed) and model (solid), the correlation of spouses' home production hours is highest when marriage market sorting is strongest (i.e. when partners' human capital is equalized $s_f \approx s_m$, around the vertical line at 'zero'). This is a natural prediction of our model: Spouses of similar human capital can better act on the hours complementarity in home production and better align their hours relative to couples with large human capital differences who tend to specialize.

Finally, households' time allocation choices in our model are also related to labor market sorting. Figure 5 (right) shows the labor market matching function *controlling* for hours worked. The difference in sorting across gender nearly vanishes both in the model (solid) and the data (dashed), relative to what we see in the left panel. In sum, the monotone equilibrium of our model—driven by home hours complementarity—fits well the rich empirical patterns of marriage sorting, labor sorting, hours allocations, and their interconnections.

Figure 5: Labor Market Matching Function, Original (left) and with Hours Partialled Out (right)

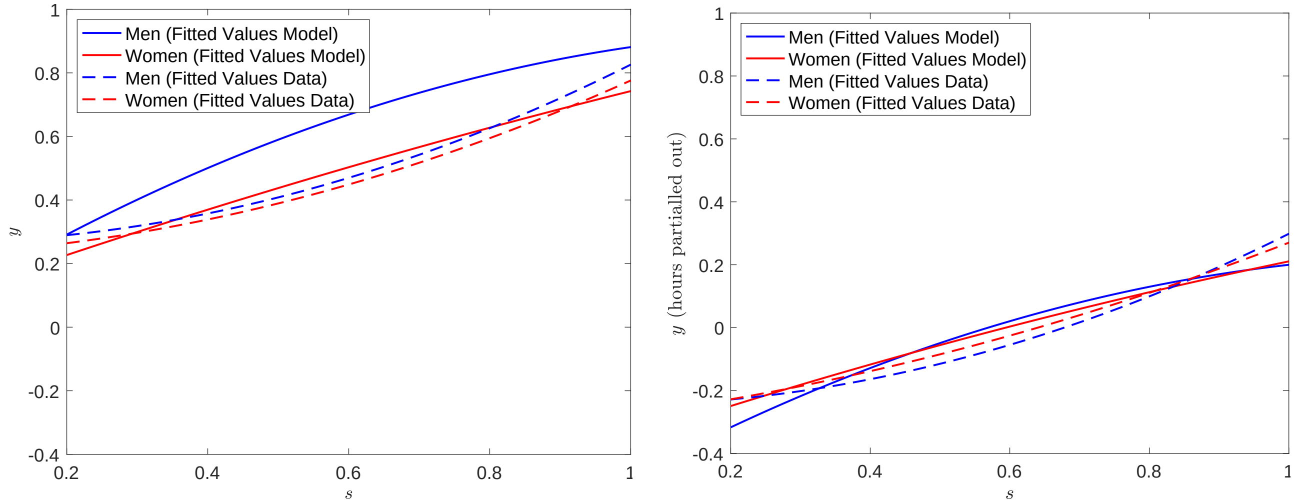

Figure 6: Labor Market Sorting and Marriage Market Sorting (left); Home Production Hours and Marriage Market Sorting (right)

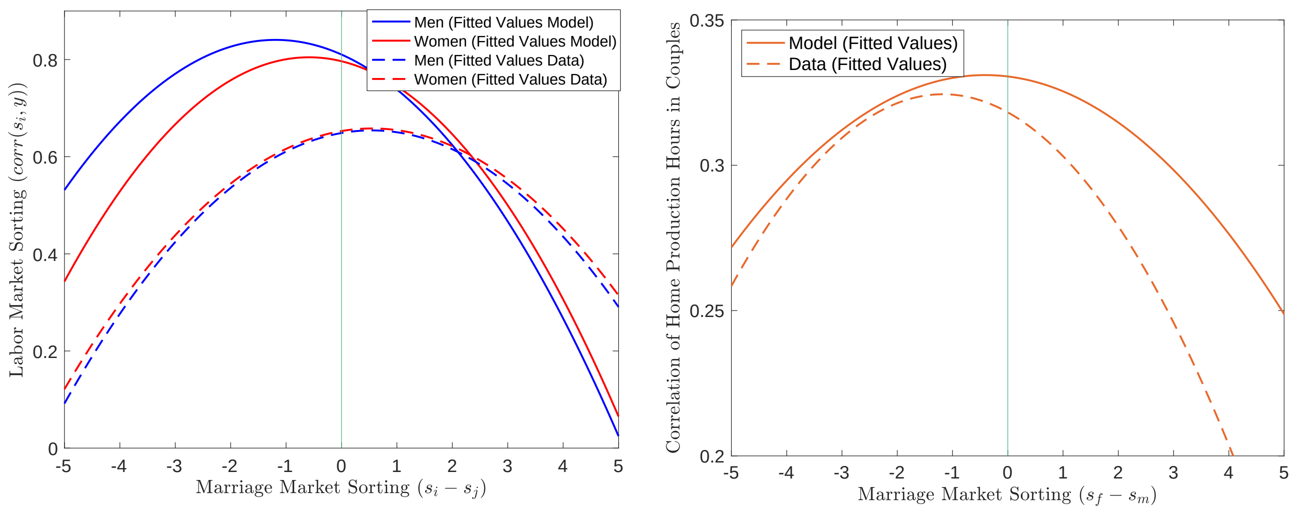

8 Application: The Drivers of Inequality

In our main quantitative exercise, we use our model to shed new light on the sources of gender disparities in the labor market and household income inequality. Our analysis focusses on two different contexts: Today's Germany (Section 8.1) and Germany over time (Section 8.2).

8.1 Inequality Through the Lens of our Model

We first focus on a recent period, 2010-2016. Throughout, we keep the focus on West Germany. We analyze the gender wage gap and income inequality within and between households through the lens of our model. We start with investigating the performance of our model in reproducing the observed inequality. We then analyze comparative statics with respect to the model's key determinants of inequality.

8.1.1 Inequality in Data and Model

To assess the extent of inequality in data and model, we focus on four measures: The gender wage gap and household income variance, including its decomposition into between and within household components. These statistics are reported in Table 3.[26] While our model underestimates the level of the income variance (83 in the model versus 98 in the data), we capture the split of within- and between-household inequality quite well (54-46 split in the model vs. 50-50 in the data). Moreover, regarding inequality within households, the model captures the gap between men and women accurately: It predicts that the share of female wages in overall household wage income is 31% (in the data, it is 33%). Last, our model produces a sizable unconditional gender wage gap (23%), slightly overestimating the observed gap (20%).

Our model is thus able to reproduce key features of observed inequality that were not targeted in estimation. This validation suggests that our model is an adequate tool through which we can investigate the main drivers of inequality, and understand the sources of changing inequality in Germany over time.

Table 3: Gender and Household Inequality

	Model	**Data**
Total Household Wage Variance	83.4822	97.7822
Within Household Wage Variance	38.0682	49.1806
... share in total variance	0.4560	0.5030
Between Household Wage Variance	45.4140	48.6017
... share in total variance	0.5440	0.4970
Share of Female Wage in Overall Household Wage Income	0.3144	0.3285
Gender Wage Gap	0.2314	0.1973

8.1.2 Comparative Statics

To highlight the key forces behind inequality we begin with comparative statics exercises in our estimated model. This will help us understand changes in inequality over time through the lens of our model below. The gender wage gap in our model is driven by (endogenous) gender differences in hours worked, and (exogenous) differences in labor productivity. In turn, differences in hours worked across gender are mostly impacted by the relative productivity of women at home, θ, the home production complementarity, ρ, and the labor market productivity wedge ψ. Clearly, if both $\theta = 0.5$ (men and women are equally productive at home) and $\psi = 1$ (men and women are equally productive in the labor market/women are not discriminated against), this would eliminate the gender wage gap. But given that $\theta \neq 0.5, \psi \neq 1$, the level of complementarities in home hours, ρ, is a third key determinant of

[26] The gender wage gap is computed as the difference in mean wages of men and women over mean men's wages. The within-component is measured by the variance of wages within a couple, averaged across all couples. The between-component is measured as the variance of the average income of each couple. Our measure of the gender wage gap includes all individuals in the sample, singles and in couples, conditional on employment. In turn, both the female's share in household income and the total income variance and its decomposition are computed based on the sample of couples. All couples are included, independent of employment status.

gender inequality in our model. We are interested in the comparative statics effects of these parameters on the gender wage gap, and also on intra and inter-household income inequality.

One can think of several policies and technological changes that impact these parameters. Anti-discrimination policies (such as gender quota or equal pay policies) can affect ψ. Childcare availability and parental leave policies (such as "daddy months") might affect θ. As the child-related tasks women were expected to perform at home are performed by someone else, gender differences in productivity of home hours are likely to decline. Further, changes in home production technologies that facilitate the house chores women traditionally specialized in (Greenwood et al., 2016; Greenwood, 2019), affect θ and ρ. Finally, an increase in the returns to investment in children (Lundberg, Pollak, and Stearns, 2016) could also impact ρ, as parents make more *joint* investments.

THE EFFECT OF ρ. We first investigate a change in home production technology that increases the complementarities in spouses' home hours. Recall that our estimate $\rho = -0.54$ indicates home hours are strategic complements. We are interested in the effects on inequality when ρ becomes even more negative, and in the underlying mechanism (changes in marriage sorting, hours, labor sorting).

Figure 7, first row, plots the effect of ρ on different inequality measures: The gender wage gap (panel a), within and between household income inequality (panel b), and overall household income inequality (panel c). It shows that a decline in ρ (moving from the right to the left on the x-axis) *decreases* the gender gap significantly. Starting from our estimate $\rho = -0.54$ and decreasing this parameter to -2 decreases the gender wage gap by almost 13%. This is due to a direct effect of complementarities on hours and several indirect effects through sorting: First, because complementarity in home production (and thus, in labor hours) among partners increases, complementarity among spouses' types in the marriage value becomes stronger, resulting in more positive assortative matching (panel d). Both, increased marriage sorting (indirectly) and stronger complementarities in home production (directly) induce spouses to better align their hours. Women increase their labor hours while men decrease theirs, leading to a smaller gender gap in labor hours (panel e), which puts downward pressure on the wage gap. Moreover, because women 'improve' a sorting-relevant attribute (work hours) relative to men, the gender gap in labor market sorting declines (panel f) reducing the gender wage gap even further.

How does this change in home production complementarities affect household income inequality? Figure 7c shows that overall income inequality *declines* with stronger complementarities. This decline is driven by the decrease in within-household inequality (mirroring the decline in the gender wage gap), which dominates the increase in between inequality that stems from stronger marriage sorting.

THE EFFECT OF θ. Next, we are interested in the effect of women's relative productivity at home on our inequality measures. The gender wage gap is increasing in θ (starting at $\theta = 0.5$), as shown in Figure 11a in Appendix D.1. Indeed, eliminating the gap in home productivity (from estimate $\theta = 0.78$ to $\theta = 0.5$) would cut the gender wage gap by almost half. The mechanism is as follows: increasing female home productivity decreases the incentive for positive marriage sorting (panel d) and pushes toward household specialization with a large gender gap in labor hours (panel e). This negatively affects women's wages directly, as well as indirectly through a larger labor market sorting gap (panel f). Interestingly, overall household income inequality increases as women become more productive at home, Figure 11c. Here, this is driven by an increase in within-household inequality (mimicking the evolution of the gender wage gap), which dominates the decline in between-household inequality driven by a drop in marriage market sorting (Figure 11b).

Figure 7: a. Gender Wage Gap, b. Household Income Variance Decomposition, c. Household Income Variance, d. Marriage Market Sorting, e. Gender Gap Labor Hours, f. Gender Gap Labor Sorting.

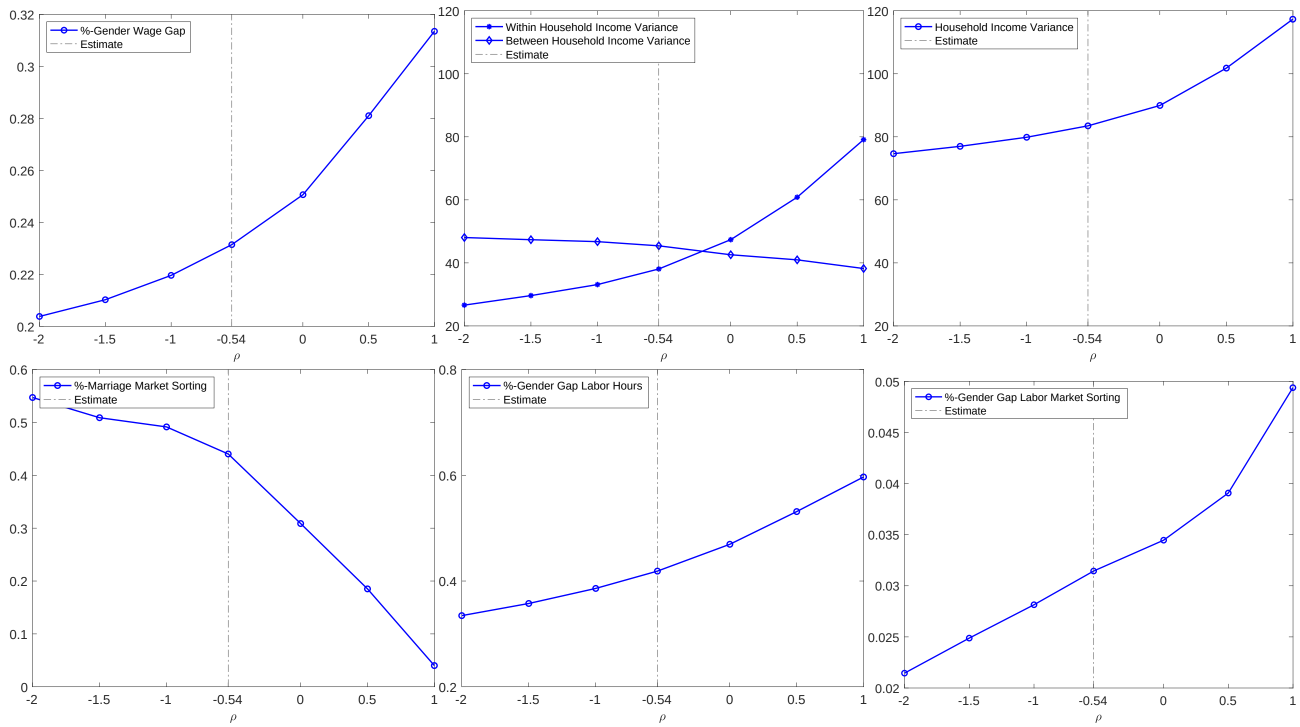

The Effect of ψ. Last, we analyze the comparative statics of the female labor market wedge. Figure 12a in Appendix D.1 shows that eliminating the wedge (increasing ψ from our estimate $\psi = 0.84$ to $\psi = 1$) would reduce the gender gap by about 25%. There is a direct positive effect of ψ on female productivity and thus wages but also several indirect effects: First, the wife's labor hours increase in productivity ψ relative to the husband's, reducing the gender hours gap (panel e) and thus gender wage gap. Second, the reduction in the gender hours gap leads to a decline in the labor market sorting gap (panel f) further curbing the gender wage gap. Third, smaller gender disparities on the labor market are associated with an increase in marriage market sorting (panel d) since in a world where men and women are more equal the motive for positive sorting strengthens. The increase in marriage sorting reinforces the drop in both hours and labor sorting gaps, further dampening the gender wage gap.

In Figure 12b and c, we study the effects of ψ on the variance of income, both within and across households. An increase in ψ leads to lower within-household inequality, but higher between-household inequality (driven by the increase in marriage sorting), with ambiguous effect on overall inequality.

Figure 8: Inequality Changes Over Time (left); Sorting and Hours Changes Over Time (right).

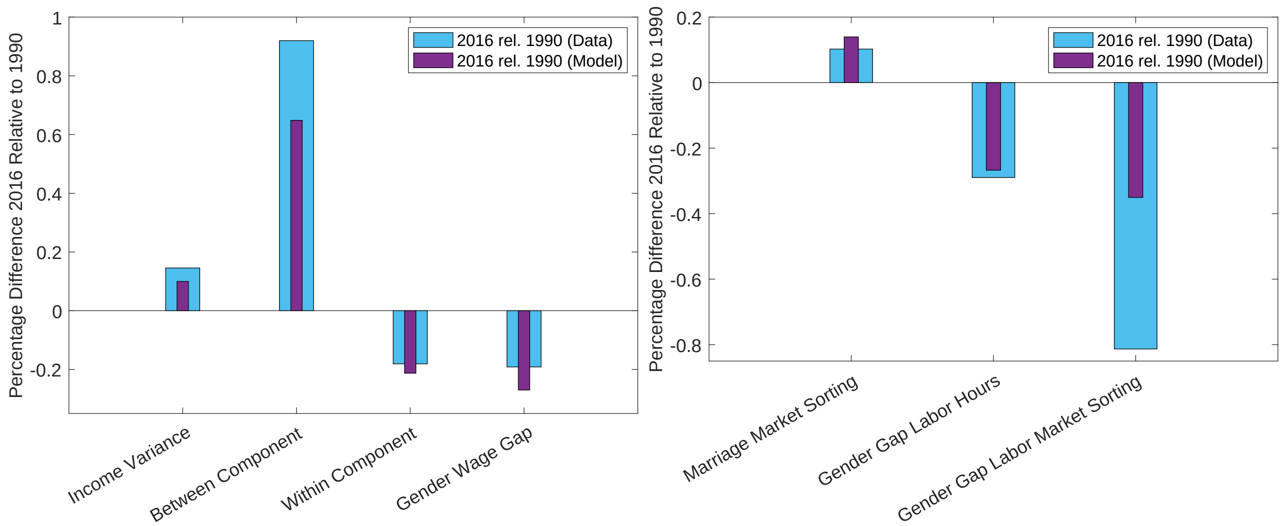

We derived several insights: First, eliminating asymmetries in productivity across gender (whether at home through $\theta \to 0.5$ or at work through $\psi \to 1$) reduces the gender wage gap. But this is not the only way to reduce gender disparities: an increase in home production complementarity (decrease in ρ, the key parameter of our model in shaping equilibrium) has qualitatively similar effects. Second, a decline in the gender wage gap tends to go hand in hand with a decline in the labor hours gap and in the labor market sorting gap and with an increase in marriage market sorting. Third, while the effect of these parameters on overall income inequality depends on the exercise, in all cases, the gender wage gap co-moves positively with within-household inequality but negatively with between-household inequality.

8.2 Inequality Over Time

Over the last few decades, inequality in Germany has changed significantly. In Figure 8, left panel, the turquoise bars show that household income variance is 15% higher today than 30 years ago, which masks diverging trends of within-household inequality (which declined by 18%) and between-household inequality (which increased by 92%). In turn, the gender wage gap declined by almost 20% over this period. At the same time, both the marriage and the labor market have undergone notable changes. The turquoise bars (right panel) show that marriage market PAM increased by 10%, while the gender gap in labor hours fell by almost 30% and the gender gap in labor market sorting, by almost 80%.

We are interested in how our model rationalizes these trends in a unified way. We first investigate how the model primitives have changed over time and how these changes affected inequality. We then ask whether the documented shifts in labor and marriage sorting amplified or mitigated inequality.

To assess over-time changes in inequality with our model, we compare our estimation from 2010-2016 with the re-estimated model in an earlier period, 1990-1996.[27] For re-estimation on the 1990-1996 sample, we re-assess the skill and job distributions for the earlier period, and re-estimate all parameters

[27] In the GSOEP, 1990 is the first year that features the time use variables used in our analysis of the later period.

except those pertaining to the labor supply preference shock, which we set to the level of our current period benchmark (Section 7.4). This is to tie our hands and force the model mechanism to explain the data, as opposed to giving changes in shock distributions a too prominent role.[28] The model fit along targeted moments is in Table 13 in Appendix D.2, which also indicates that both labor and marriage market underwent statistically significant changes over time (column 5). Regarding the un-targeted inequality moments of the data (Figure 8, left panel), the model replicates the over-time changes quite well, where the turquoise bars indicate changes in the data and the purple bars, changes in the model.

To understand the driving forces behind the inequality changes, we now zoom further into the model. We compare the parameter estimates for both periods in Table 14, Appendix D.2. There have been significant changes in home production with today's Germany being characterized by a lower ρ (drop from -0.16 to -0.54, indicating increased complementarity in spouses' home hours) and a lower θ (drop from 0.88 to 0.78, meaning men became relatively more productive at home over time); and a narrowing labor productivity wedge ψ (increase from 0.76 to 0.84, raising relative female productivity). These changes indicate that Germany has become an economy with more gender equality both at home and at work. In turn, the labor market technology has become more convex in effective skills, resembling skill-biased technological change, and it has a higher TFP than before.

How much of the documented changes in inequality can be explained by these changes of model parameters? Figure 9 provides a detailed decomposition. The purple bars re-display the overall change in inequality produced by the model—where we account for *all* parameters changes over time. The remaining bars give the percentage change in inequality outcomes between 1990-1996 and 2010-2016 if one parameter group changes *in isolation* while the others remain fixed at the 1990-1996 level: We consider changes in the labor market production function (blue), home production (orange), labor productivity wedge (yellow), and human capital distribution (green).

In line with our comparative statics exercises, the documented changes in home production technology reduced gender disparities (gender wage gap and within-household inequality) as well as overall household inequality, while they fueled between-household inequality. Figure 9 (orange bars) shows that these effects are also quantitatively sizable. If only home production had changed over time, within-household inequality would have declined by 30% (accounting for more than the observed change) and the gender wage gap by 14% (accounting for more than 70% of the observed drop). Home production changes were thus the biggest driver behind the decline in gender inequality. In turn, home production shifts put upward pressure on between-household inequality, accounting for almost 20% of the observed increase. But since this effect was dominated by the downward pressure on within-inequality, the net effect of technological change in home production on overall household income inequality was negative. Splitting home production further into the contributions of our model's key parameters θ and ρ (Figure 13 in Appendix D.2) reveals that changes in relative productivity parameter θ were the main driver behind the inequality shifts (accounting for around 2/3 of the total home production effects), while the

[28] We did have to free up the scale of marriage shocks in 1990-1996 in order to give the model a chance to match the data.

impact of complementarity parameter ρ was smaller but still sizable (around 1/3 of the effects). The effects of changes in the labor market wedge ψ on inequality (yellow bars, Figure 9)—while qualitatively similar to those of home production technology—were quantitatively smaller. Finally, changes in labor market technology (blue bars) fueled inequality across the board, significantly pushing up household income variance (through both between- and within-components) and preventing gender inequality from falling even further. Thus, technological change in home production and in labor market production have pushed inequality, and especially gender disparities, in opposite directions.

Figure 9: Mechanism Behind Inequality Changes

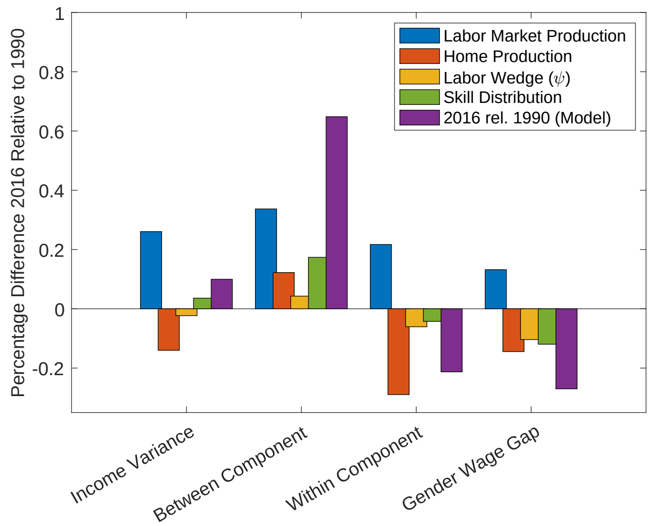

Our comparative statics in Section 8.1.2 clarify the mechanism of why the estimated changes of home production technology and the labor wedge push towards more gender equality. Both changes induced women to work more (leading to a decline in the gender gap of labor hours), which in turn made women sort relatively better on the labor market (reducing the gender gap in labor sorting). More gender parity in labor market outcomes in turn strengthened the desire for positive sorting in marriage, reinforcing the push towards more equal labor (and home) hours across gender. Figure 8, right panel, demonstrates that these shifts were not only present in the model (purple bars) but also in the data (turquoise bars). Our evidence and estimates suggest that Germany underwent significant changes over the last decades towards an equilibrium that resembles the monotone equilibrium from our theory, with stronger home production complementarities and, consequently, increased marriage sorting as well as stronger co-movements of spouses' hours, labor market sorting, and wages.

We end by returning to a key feature of our model: equilibrium sorting in both labor and marriage markets. We assess the quantitative role of changes in marriage and labor market sorting for inequality shifts. Between 1990-96 and 2010-16, positive marriage sorting has increased by around 10% and positive labor sorting, by 8%. We compute the elasticity of each inequality outcome with respect to sorting in each market as (%ΔInequality)/(%ΔLabor Sorting) and (%ΔInequality)/(%ΔMarriage Sorting). *Inequality* refers to one of our four inequality outcomes (gender wage gap, household income variance, within/between component) and the percentage change is computed between the baseline model in 2010-2016 and the counterfactual model. This counterfactual inputs the estimated parameters from 2010-2016 but keeps either labor market or marriage market sorting constant at the past period's (1990-1996) level.[29] This way, we isolate the role of the observed changes in sorting for inequality shifts.

Table 4: Elasticity of Inequality with Respect to Sorting

	Gender Wage Gap	**Income Variance**	**Within Variance**	**Between Variance**
Marriage Market Sorting	-0.0036	0.0117	-0.1173	0.1234
Labor Market Sorting	-1.2349	0.0227	-0.7356	0.7351

Table 4 reports the elasticities. We find that *both* marriage and labor sorting have had mitigating impact on gender inequalities (wage gap and within-household inequality) and have amplified overall inequality and between-household inequality. For instance, a 1% increase in marriage sorting has decreased within-household inequality by 0.117%, while it increased between-household inequality by 0.123%. The elasticity of the gender wage gap is also negative, albeit smaller. Stronger marriage market sorting generated more balanced labor market outcomes—in hours, sorting, and pay—across gender.

The effects of changes in labor sorting on inequality are even larger. A 1% rise in labor market sorting has increased the between-household income variance by 0.735%. In turn, a 1 % increase in labor sorting has reduced the gender wage gap by 1.235% and within inequality by 0.736%. Surprisingly at first sight, the increase in labor sorting over the past decades significantly *narrowed* gender disparities. The reason is that this increase was predominantly driven by women's improved labor sorting (the gender gap in labor sorting has declined over time, Figure 8, right panel), helping them to catch up with men's pay. Stronger positive sorting between workers and jobs—when over-proportionally benefitting women—can spur gender convergence in labor market outcomes.

9 Conclusion

Employers value workers not only for their skills but also for their time input. In such a setting, if labor supply decisions are made at the household level so that they depend on the characteristics of *both* spouses, then marriage market sorting affects labor market sorting. In turn, if individuals anticipate their hours choices as well as their sorting and pay in the labor market when deciding whom to marry, then labor market sorting affects marriage market sorting. The interaction of *both* the marriage and the labor market crucially impacts inequality across gender and within/between households. And policies affecting who marries whom (such as tax policies) or home production technology (such as parental leave or universal childcare) can therefore mitigate or amplify inequality, calling for a better understanding of these spillovers across markets.

[29] To implement the past marriage sorting in the counterfactual model, we adjust σ_β^M. In turn, to implement the past labor sorting, we take into account labor market matching $\mu(\tilde{s})$ from the past period when computing the wage function.

The interplay between labor market and marriage market and its effect on inequality are at the center of this paper. We build a novel equilibrium model in which households' labor supply choices form the natural link between the two markets and their sorting margins. We first show that in theory, the nature of home production—whether partners' hours are complements or substitutes—shapes marriage market sorting, labor supply choices, and labor market sorting in equilibrium.

We then ask what is the nature of home production in the data. To this end, we estimate our model on data from today's Germany and find that spouses' home hours are strategic *complements*, pushing towards positive sorting in both markets and co-movement of labor hours of spouses. This is in contrast to what would happen in a 'traditional' economy based on substitution in home production and specialization of spouses. Investigating the key drivers behind inequality based on primitives, we find that the gender wage gap and within-household income inequality would decrease not only if gender productivity differences at home or in the labor market were reduced, but also if home production hours were even more complementary among partners. Home production complementarities induce spouses to split their time similarly between work in the market and at home. And they also increase marriage sorting and reduce the gender gap in labor sorting, both mitigating gender disparities further.

Our main quantitative exercise analyzes how our model can rationalize changes in inequality over time. We find that home production hours of spouses have become more complementary over time and that this technological change in home production can account for a significant part of the decline in gender inequality in Germany. In contrast, technological change in the labor market has fueled inequality across the board, including gender gaps. Highlighting the unique feature of our model, we show that sorting on both markets has significant quantitative effects on inequality: We find that *both* stronger marriage market sorting and labor market sorting over time have amplified overall inequality and between-household inequality, but have had a mitigating impact on gender inequalities (wage gap and within-household inequality)—highlighting a new role of sorting for gender convergence in pay.

References

D. Aaronson and E. French. The effect of part-time work on wages: Evidence from the social security rules. *Journal of Labor Economics*, 22(2):329–252, 2004.

J. Adda, C. Dustmann, and K. Stevens. The career costs of children. *Journal of Political Economy*, 125 (2):293–337, 2017.

S. Albanesi and C. Olivetti. Home production, market production and the gender wage gap: Incentives and expectations. *Review of Economic Dynamics*, 12(1):80 – 107, 2009.

I. Andrews, M. Gentzkow, and J. M. Shapiro. Measuring the sensitivity of parameter estimates to estimation moments. *The Quarterly Journal of Economics*, 132(4):1553–1592, 2017.

N. Angelov, P. Johansson, and E. Lindahl. Parenthood and the gender gap in pay. *Journal of Labor Economics*, 34(3):545–579, 2016.

G. Anyaegbu. Using the OECD equivalence scale in taxes and benefits analysis. *The Labour Gazette*, 4 (1):49–54, 2010.

J. Bagger and R. Lentz. An Empirical Model of Wage Dispersion with Sorting. *The Review of Economic Studies*, 86(1):153–190, 2018.

G. S. Becker. A theory of marriage: Part I. *Journal of Political Economy*, 81(4):813–846, 1973.

M. Bertrand, C. Goldin, and L. F. Katz. Dynamics of the gender gap for young professionals in the financial and corporate sectors. *American Economic Journal: Applied Economics*, 2(3):228–55, 2010.

A. Bick, A. Blandin, and R. Rogerson. Hours and wages. 2020.

S. Bonhomme, T. Lamadon, and E. Manresa. A distributional framework for matched employer employee data. *Econometrica*, 87(3):699–739, 2019.

J. N. Brown and H. S. Rosen. On the Estimation of Structural Hedonic Price Models. *Econometrica*, 50(3):765–68, May 1982.

M. Browning, P.-A. Chiappori, and Y. Weiss. *Economics of the Family*. Cambridge Books. Cambridge University Press, 2014.

D. Card, J. Heining, and P. Kline. Workplace Heterogeneity and the Rise of West German Wage Inequality. *The Quarterly Journal of Economics*, 128(3):967–1015, 2013.

P. Chiappori, M. Iyigun, and Y. Weiss. Investment in schooling and the marriage market. *The American Economic Review*, 99(5):1689–1713, 2009.

P.-A. Chiappori, B. Salanié, and Y. Weiss. Partner choice, investment in children, and the marital college premium. *American Economic Review*, 107(8):2109–67, August 2017. doi: 10.1257/aer.20150154. URL https://www.aeaweb.org/articles?id=10.1257/aer.20150154.

P.-A. Chiappori, M. Costa-Dias, and C. Meghir. The marriage market, labor supply, and education choice. *Journal of Political Economy*, 126(S1):S26–S72, 2018.

P. A. Chiappori, M. C. Dias, and C. Meghir. Changes in assortative matching: Theory and evidence for the us and the uk. 2020.

H. L. Cole, G. J. Mailath, and A. Postlewaite. Efficient non-contractible investments in large economies. *Journal of Economic Theory*, 101(2):333 – 373, 2001.

P. Cortés and J. Pan. When time binds: Substitutes for household production, returns to working long hours, and the skilled gender wage gap. *Journal of Labor Economics*, 37(2):351–398, 2019.

G. Cubas, C. Juhn, and P. Silos. Coordinated work schedules and the gender wage gap. Working Paper, 2020.

M. C. Dias, R. Joyce, and F. Parodi. The gender pay gap in the uk: Children and experience in work. IFS Working Papers W18/02, 2018.

L. Eika, M. Mogstad, and B. Zafar. Educational assortative mating and household income inequality. *Journal of Political Economy*, 127(6):2795–2835, 2019.

I. Ekeland, J. J. Heckman, and L. Nesheim. Identification and estimation of hedonic models. *Journal of Political Economy*, 112(S1):S60–S109, February 2004.

A. Erosa, L. Fuster, G. Kambourov, and R. Rogerson. Hours, occupations, and gender differences in labor market outcomes. Working Paper 23636, National Bureau of Economic Research, 2017.

R. Fernández, N. Guner, and J. Knowles. Love and Money: A Theoretical and Empirical Analysis of Household Sorting and Inequality. *The Quarterly Journal of Economics*, 120(1):273–344, 2005.

L. Flabbi, C. J. Flinn, and M. Salazar-Saenz. Simultaneous search in the labor and marriage markets with endogenous schooling decisions. 2020.

G.-L. Gayle and L. Golan. Estimating a Dynamic Adverse-Selection Model: Labour-Force Experience and the Changing Gender Earnings Gap 1968–1997. *The Review of Economic Studies*, 79(1):227–267, 2011.

G.-L. Gayle and A. Shephard. Optimal taxation, marriage, home production, and family labor supply. *Econometrica*, 87(1):291–326, 2019.

D. Gicheva. Working long hours and early career outcomes in the high-end labor market. *Journal of Labor Economics*, 31(4):785–824, 2013.

C. Goldin. A grand gender convergence: Its last chapter. *American Economic Review*, 104(4):1091–1119, April 2014.

M. Goussé, N. Jacquemet, and J.-M. Robin. Household labour supply and the marriage market in the uk, 1991-2008. *Labour Economics*, 46:131 – 149, 2017a.

M. Goussé, N. Jacquemet, and J.-M. Robin. Marriage, labor supply, and home production. *Econometrica*, 85(6):1873–1919, 2017b.

J. Greenwood. *Evolving Households The Imprint of Technology on Life*. MIT Press, 2019.

J. Greenwood, N. Guner, G. Kocharkov, and C. Santos. 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, 2016.

J. Greenwood, N. Guner, and G. Vandenbroucke. Family economics writ large. *Journal of Economic Literature*, 55(4):1346–1434, December 2017.

M. Hagedorn, T. H. Law, and I. Manovskii. Identifying equilibrium models of labor market sorting. *Econometrica*, 85(1):29–65, 2017.

J. J. Heckman. Sample selection bias as a specification error. *Econometrica*, 47(1):153–161, 1979.

G. L. Hunt. Alternative nested logit model structures and the special case of partial degeneracy. *Journal of Regional Science*, 40(1):89–113, 2000.

M. Keane, P. Todd, and K. I. Wolpin. The structural estimation of behavioral models: Discrete choice dynamic programming methods and applications. volume 4A, chapter 04, pages 331–461. Elsevier, 1 edition, 2011.

H. Kleven, C. Landais, and J. E. Søgaard. Children and gender inequality: Evidence from denmark. *American Economic Journal: Applied Economics*, 11(4), 2019.

P. Legros and A. F. Newman. Beauty Is a Beast, Frog Is a Prince: Assortative Matching with Nontransferabilities. *Econometrica*, 75(4):1073–1102, 07 2007.

I. Lindenlaub. Sorting Multidimensional Types: Theory and Application. *The Review of Economic Studies*, 84(2):718–789, 2017.

I. Lindenlaub and F. Postel-Vinay. Multidimensional sorting under random search. 2020.

J. Lise and F. Postel-Vinay. Multidimensional skills, sorting, and human capital accumulation. *American Economic Review*, Forthcoming.

J. Lise and S. Seitz. Consumption Inequality and Intra-household Allocations. *The Review of Economic Studies*, 78(1):328–355, 01 2011.

S. Lundberg, R. Pollak, and J. Stearns. Family inequality: Diverging patterns in marriage, cohabitation, and childbearing. 2016.

A. Mas and A. Pallais. Valuing alternative work arrangements. *American Economic Review*, 107(12): 3722–59, 2017.

M. Mazzocco. Household intertemporal behaviour: A collective characterization and a test of commitment. *The Review of Economic Studies*, 74(3):857–895, 2007.

D. McFadden. A method of simulated moments for estimation of discrete response models without numerical integration. *Econometrica*, 57(5):995–1026, 1989.

OECD. Starting strong iv. early childhood education and care data country note. germany. Technical report, OECD Publishing, Paris, 2016.

A. Pakes and D. Pollard. Simulation and the asymptotics of optimization estimators. *Econometrica*, 57 (5):1027–1057, 1989.

L. Pilossoph and S. L. Wee. Assortative matching and income inequality: A structural approach. 2019a.

L. Pilossoph and S. L. Wee. Household search and the marital wage premium. 2019b.

A. Reynoso. The impact of divorce laws on the equilibrium in the marriage market. 2019.

S. Rosen. Hedonic prices and implicit markets: Product differentiation in pure competition. *Journal of Political Economy*, 82(1):34–55, Jan.-Feb. 1974.

L. S. Shapley and M. Shubik. The assignment game I: The core. *International Journal of Game Theory*, 1(1):111–130, 1971.

A. Voena. Yours, mine, and ours: Do divorce laws affect the intertemporal behavior of married couples? *American Economic Review*, 105(8):2295–2332, 2015.

Appendix

A Theory

A.1 Monotone Equilibrium: Derivations and Proofs

A.1.1 Stability

We begin by analyzing stability of equilibrium. In particular, since our comparative statics apply to any stable equilibrium, where we refer to stability of the household problem. (In turn, stability in the marriage and labor market is trivially satisfied in the competitive equilibrium.) We therefore first define stability of the household problem.

Definition 3. *The equilibrium in the household stage, given by* (h_f, h_m)*, is stable for a given wage function* w *if*

$$w''(\tilde{x}_m)e_h(\tilde{x}_m, h_m)^2 + w'(\tilde{x}_m)e_{hh}(\tilde{x}_m, h_m) + 2p_{\ell_m\ell_m}(1-h_m, 1-h_f) < 0 \tag{23}$$

$$w''(\tilde{x}_f)e_h(\tilde{x}_f, h_f)^2 + w'(\tilde{x}_f)e_{hh}(\tilde{x}_f, h_f) + 2p_{\ell_f\ell_f}(1-h_m, 1-h_f) < 0 \tag{24}$$

$$\begin{aligned}&\left(w''(\tilde{x}_m)e_h(\tilde{x}_m, h_m)^2 + w'(\tilde{x}_m)e_{hh}(\tilde{x}_m, h_m) + 2p_{\ell_m\ell_m}(1-h_m, 1-h_f)\right)\\ &\times\Big(w''(\tilde{x}_f)e_h(\tilde{x}_f, h_f)^2 + w'(\tilde{x}_f)e_{hh}(\tilde{x}_f, h_f) + 2p_{\ell_f\ell_f}(1-h_m, 1-h_f)\Big)\\ &\qquad -4p_{\ell_m\ell_f}(1-h_m, 1-h_f)p_{21}(1-h_m, 1-h_f) > 0.\end{aligned} \tag{25}$$

In what follows, we will omit the arguments of the functions to make the notation more concise.

We now explain what is behind this definition of stability. Stability is reflected in the properties of the spouses' 'best response' functions, which are implicitly given by the household's FOCs (8)-(9). The first FOC gives rise to a 'best response' function r_f for the woman's hours to man's hours, given by $h_f^* = r_f(h_m^*)$ for all $h_m^* \in [0,1]$. The second FOC gives rise to a 'best response' function r_m for the man's hours to the woman's hours, given by $h_m^* = r_m(h_f^*)$ for all $h_f^* \in [0,1]$. An equilibrium in the household stage is any solution (h_m^*, h_f^*) to (8)–(9), an intersection of the best response functions. It is stable if small perturbations in the hours choices induce the agents to converge back to equilibrium. This is true iff (23)–(25) hold:

Inequalities (23)–(24) require that, given the labor hour choice of the partner, own hours adjust properly if the marginal benefit from 'investing' (i.e. the marginal wage benefit from more labor hours) is higher/lower than the marginal cost (forgone home production), i.e., if higher then hours go up; if lower then hours go down.

In turn, inequality (25) requires that, at the crossing point of the two best response functions, $|\partial r_f/\partial h_m^*| < |\partial r_m^{-1}/\partial h_m^*|$. That is, the slope of r_f is smaller than the slope of the inverse of r_m when they cross in the (h_m^*, h_f^*)-space, or r_m crosses r_f from below (above) when the BR-functions are upward (downward) sloping at the crossing. To see that inequality (25) requires that, at the crossing of the two best response functions, $|\partial r_f/\partial h_m^*| \leq |\partial r_m^{-1}/\partial h_m^*|$, note that $|\partial r_m^{-1}/\partial h_m^*| > |\partial r_f/\partial h_m^*|$ is equivalent to

$$\begin{aligned}
\Leftrightarrow \quad & \left| - \frac{w''e_h^2 + w'e_{hh} + 2p_{\ell_f\ell_f}}{2p_{\ell_m\ell_f}} \right| \left| - \frac{w''e_h^2 + w'e_{hh} + 2p_{\ell_m\ell_m}}{2p_{\ell_m\ell_f}} \right| && > 1 \\
\Leftrightarrow \quad & \left| - \left(w''e_h^2 + w'e_{hh} + 2p_{\ell_f\ell_f} \right) \right| \left| - \left(w''e_h^2 + w'e_{hh} + 2p_{\ell_m\ell_m} \right) \right| && > (2p_{\ell_m\ell_f})^2 \\
\Leftrightarrow \quad & \left(w''e_h^2 + w'e_{hh} + 2p_{\ell_f\ell_f} \right) \left(w''e_h^2 + w'e_{hh} + 2p_{\ell_m\ell_m} \right) && > (2p_{\ell_m\ell_f})^2
\end{aligned}$$

where the last inequality follows from conditions (23)–(24), ensuring that in a stable equilibrium, both terms in brackets on the LHS are negative. The crossing of the best response functions described by (25) guarantees that small perturbations away from the equilibrium hours induce dynamics so that the resulting hours adjustments make the household converge back to the equilibrium hours.

With Definition 3 at hand, we can now show the following.

Lemma 1. *For any wage function w with the property that $w_{h_i} = w'(\tilde{x}_i)e_h(x_i, h_i)$ is strictly positive and finite for all $h_i \in [0,1]$, a stable equilibrium of the household problem exists.*

Proof. We first show that, given $h_m^* \in [0,1]$, there exists a solution to (8), $w'e_h = 2p_{\ell_f}$ that satisfies also (24). To see this, notice that, for any h_m^*, the LHS, $w_{h_f} = w'e_h$, is always strictly positive and finite under the premise, while the RHS goes to infinity as h_f^* goes to one (i.e. when $1 - h_f$ goes to zero), and to zero as h_f^* goes to zero (and $1 - h_f$ to one) by the assumed Inada conditions on p. Hence, by the Intermediate Value Theorem, there is at least one solution to (8), and the first solution satisfies (23) (the RHS crosses the LHS from below when plotted against h_f). Denote this solution by $h_f^* = r_f(h_m^*)$ for each h_m^*, and notice that it is continuously differentiable in h_m^*.

Similarly, given $h_f^* \in [0,1]$, there exists a solution to (9) that satisfies also (23). We denote this solution by $h_m^* = r_m(h_f^*)$ for each h_f^*, which is continuously differentiable in h_f^*.

Next, we note that $r_f(0) > 0$ and $r_f(1) < 1$, which follows from the strictly positive and finite value of the LHS in (8), $w'e_h = 2p_{\ell_f}$, and the boundary properties of p_{ℓ_f}. Similarly, $r_m(0) > 0$ and $r_m(1) < 1$. Since r_m and r_f are continuous functions, there exists a pair (h_f^*, h_m^*) such that $h_f^* = r_f(h_m^*)$ and $h_m^* = r_m(h_f^*)$ and that also satisfies (25), i.e. $|\partial r_f/\partial h_m^*| \leq |\partial r_m^{-1}/\partial h_m^*|$.

Thus, a stable equilibrium of the household problem exists.[30] $\square$

Lemma 1 shows that for *any* given wage function that satisfies the stated assumption on w_{h_i}, there exists at least one stable equilibrium. We now argue that, as a result, also for the *equilibrium* wage function, there exists at least one stable equilibrium. To show this, the equilibrium wage function needs to satisfy the property that w_{h_i} is strictly positive and finite. It is given by (2), and so $w_{h_i} = z_{\tilde{x}}(\tilde{x}_i, \mu(\tilde{x}_i))e_h(x_i, h_i)$, which under the assumption of the model is strictly positive and finite.

[30] To see this graphically in the (h_m^*, h_f^*) space, notice that r_f starts above the inverse of r_h and ends below. Continuity implies there is a crossing between the two, and the first one is such that the inverse of r_m crosses the r_f from below. This implies that at the crossing point $|\partial r_f/\partial h_m^*||\partial r_m/\partial h_f^*| \leq 1$.

Thus, whenever an equilibrium exists, we know that there exists at least one stable one. And any stable equilibrium satisfies the conditions of Definition 3, which we use to sign our comparative statics of the household problem.

A.1.2 Spouses' Hours as Strategic Complements or Substitutes

Totally differentiating each FOC w.r.t. to (h_f, h_m), we obtain the slopes of the 'best response' functions for women (men) to men's (women's) hours, respectively:

$$
\begin{aligned}
0 &= (w'' e_h^2 + w' e_{hh}) dh_f + 2p_{\ell_m \ell_f} dh_m + 2p_{\ell_f \ell_f} dh_f \\
\Leftrightarrow \quad \frac{dh_f}{dh_m} &= -\frac{2p_{\ell_m \ell_f}}{w'' e_h^2 + w' e_{hh} + 2p_{\ell_f \ell_f}} \qquad (26) \\
0 &= (w'' e_h^2 + w' e_{hh}) dh_m + 2p_{\ell_m \ell_f} dh_f + 2p_{\ell_m \ell_m} dh_m \\
\Leftrightarrow \quad \frac{dh_m}{dh_f} &= -\frac{2p_{\ell_m \ell_f}}{w'' e_h^2 + w' e_{hh} + 2p_{\ell_m \ell_m}} \qquad (27)
\end{aligned}
$$

In any stable equilibrium (since the denominators of these expressions are negative), the best response functions are upward (downward) sloping, and thus hours are strategic complements (substitutes), if p is supermodular (submodular).

A.1.3 Proof of Proposition 1

We now check that under the specified conditions the equilibrium is monotone as specified in Definition 2.

1. PAM in the labor market in $(y, \tilde{x}_i)$ materializes due to $z_{12} > 0$.

2. The properties of how labor hours depend on own type and partner type follow from the system of FOCs of the household problem. We differentiate system (8) - (9) w.r.t. (h_f, h_m, x_f), taking as given the equilibrium wage function and the equilibrium marriage market matching function:

$$
\begin{aligned}
A\frac{\partial h_f}{\partial x_f} + B\frac{\partial h_m}{\partial x_f} &= -(w'' e_{x_f} e_h + w' e_{x_f h}) \\
C\frac{\partial h_f}{\partial x_f} + D\frac{\partial h_m}{\partial x_f} &= -\eta'(w'' e_{x_m} e_h + w' e_{x_m h})
\end{aligned}
$$

where

$$
\begin{aligned}
A &:= w'' e_h^2 + w' e_{hh} + 2p_{\ell_f \ell_f} \\
B &:= 2p_{\ell_m \ell_f} \\
C &:= 2p_{\ell_m \ell_f} \\
D &:= w'' e_h^2 + w' e_{hh} + 2p_{\ell_m \ell_m}
\end{aligned}
$$

Denote $|H| := AD - BC$, which is the determinant of the Hessian of the household problem and positive in any *stable* equilibrium, see stability condition (25) of Definition 3. Solving the system yields

$$\begin{aligned}
\frac{\partial h_f}{\partial x_f} &= \frac{-\left(w'' e_{x_f} e_h + w' e_{x_f h}\right)\left(w'' e_h^2 + w' e_{hh} + 2p_{\ell_m \ell_m}\right) + 2p_{\ell_m \ell_f}\left(w'' e_h e_{x_m} + w' e_{x_m h}\right)\eta'}{|H|} \\
&= \frac{-(u)_{h_m h_m}(w)_{x_f h} + 2p_{\ell_m \ell_f}(w)_{x_m h}\eta'}{|H|} \\
\frac{\partial h_m}{\partial x_f} &= \frac{-\left(w'' e_h^2 + w' e_{hh} + 2p_{\ell_f \ell_f}\right)(w'' e_h e_{x_m} + w' e_{x_m h})\eta' + 2p_{\ell_m \ell_f}(w'' e_{x_f} e_h + w' e_{x_f h})}{|H|} \\
&= \frac{-(u)_{h_f h_f}(w)_{x_m h}\eta' + 2p_{\ell_m \ell_f}(w)_{x_f h}}{|H|}.
\end{aligned}$$

where $(u)_{h_m h_m} = w'' e_h^2 + w' e_{hh} + 2p_{\ell_m \ell_m}$ and $(u)_{h_f h_f} = w'' e_h^2 + w' e_{hh} + 2p_{\ell_f \ell_f}$ are derived from $u(w(\tilde{x}_f) + w(\tilde{x}_m) - \overline{v} + p, p) = w(\tilde{x}_m) + w(\tilde{x}_f) - \overline{v} + 2p(1 - h_m, 1 - h_f)$, see (7), and where we again denote by $h = h_i(\cdot, \cdot)$ the hours argument in e. These expressions are equivalent to (10) and (11). They are positive in any stable equilibrium (where conditions (23)-(25) from Definition 3 are satisfied), meaning hours are increasing in own types and in partner's types, given that (i) hours are complementary in home production (p supermodular so that $p_{\ell_m \ell_f} > 0$), (ii) matching on the marriage market is PAM ($\eta' > 0$), for which we will provide conditions below, and (iii) wages are convex in effective types (yielding $(w)_{x_f h} > 0$ and $(w)_{x_m h} > 0$), which can be ensured from primitives if z is weakly convex. To see this note that:

$$\begin{aligned}
(w)_{x_f h} &= w' e_{x_f h} + w'' e_{x_f} e_h \\
&= z_{\tilde{x}} e_{x_f h} + (z_{\tilde{x}\tilde{x}} + z_{\tilde{x}y}\mu') e_{x_f} e_h,
\end{aligned}$$

and similarly for $(w)_{x_m h}$. Analogously, we can compute how hours respond to changes in male types (12) and (13) (omitted here for brevity).

3. In a stable equilibrium, PAM in the marriage market results if (as indicated in the text)

$$\begin{aligned}
\Phi_{x_m x_f} &= (w)_{x_f h} \frac{\partial \hbar_f}{\partial x_m} \\
&= 2\frac{p_{\ell_m \ell_f}(w)_{x_f h}(w)_{x_m h}}{|H|} > 0
\end{aligned}$$

where $\hbar_i := h_i(x_m, x_f)$ and where we substituted into the first line the 'partial' equilibrium comparative static $\partial \hbar_f / \partial x_m$ (see below) to obtain to the second line. Thus, $\Phi_{x_m x_f} > 0$ if $z_{\tilde{x}\tilde{x}} \geq 0$ (which renders $(w)_{x_i h} > 0$, see part 2. above) and $p_{\ell_m \ell_f} > 0$.

To show that $\partial \hbar_f / \partial x_m > 0$ (i.e. hours are increasing in partner's type for *any* couple (x_f, x_m), not only along the equilibrium assignment, which is to be solved for in this step), we differentiate system

(8) - (9) w.r.t. x_m for any given x_f, not taking η into account:

$$\frac{\partial \hbar_m}{\partial x_m} = \frac{-(w)_{x_m h}(u)_{h_f h_f}}{|H|}$$
$$\frac{\partial \hbar_f}{\partial x_m} = \frac{2(w)_{x_m h} p_{\ell_m \ell_f}}{|H|}$$

where we used the second expression to sign $\Phi_{x_m x_f}$ above.

A.1.4 Proof of Proposition 2

We check that the properties of non-monotone equilibrium are satisfied under the specified conditions.

1. PAM in the labor market in $(y, \tilde{x}_i)$ materializes due to $z_{\tilde{x}y} > 0$.

2. The properties of how labor hours depend on own type and partner type follow straight from expressions (10)–(13): $\partial h_f / \partial x_f > 0$ in a stable equilibrium if $p_{\ell_m \ell_f} < 0$ and $\eta' < 0$ as well as $z_{\tilde{x}\tilde{x}} > 0$ (which renders $(w)_{x_i h} > 0$). Further, $\partial h_f / \partial x_m < 0$ under the same conditions. It remains to verify that $\eta' < 0$, see 3 below.

3. NAM in the marriage market results if

$$\Phi_{x_m x_f} = 2 \frac{p_{\ell_m \ell_f}(w)_{x_f h}(w)_{x_m h}}{|H|} < 0$$

Thus, in a stable eq., $\Phi_{x_m x_f} < 0$ if $z_{\tilde{x}\tilde{x}} > 0$ (which renders $(w)_{x_i h} > 0$ for $i \in \{f, m\}$) and $p_{\ell_m \ell_f} < 0$.

B Solution of the Quantitative Model

The solution of our quantitative model consists of solving for a fixed point in the wage function (as a function of effective types) such that under this wage function, marital choices, household labor supply, and labor market sorting are all consistent. That is, we find the market-clearing wage function that induces households that form in the marriage market to optimally supply labor (pinning down their effective types) such that, when optimally sorting into firms on the labor market, this gives rise to that exact same wage function.

We first solve for the optimal matching in the marriage market and households' labor supply choices *given* a wage function. Given the induced labor supply decisions, individuals optimally match with firms on the labor market. Sorting in the labor market endogenously determines a new wage function (again as a function of effective types) that supports this particular matching. Given this new wage function, new marriage and labor supply decisions are made that, in turn, again affect wages in the labor market. We iterate between the problem of households on the one hand and that of workers and firms on the other until the wage function converges (until a fixed point in the wage function is found).

We next describe the solution in each decision stage, starting backwards from the labor market and then going to household and marriage problems. Finally, we outline the algorithm to find the fixed point.

B.1 Partial Equilibrium in the Labor Market ([lpe])

First, we show how we solve for the matching and wage functions in the labor market, (μ, w). Consider our *exogenous* distribution of firms, $y \sim G$, and *any given* distribution of effective types, $\tilde{s} \sim \tilde{N}_s$. Note that even though $\tilde{N}_s$ is an endogenous object in our model, from a partial equilibrium perspective where marital and household choices are taken as given, firms take the distribution $\tilde{N}_s$ as fixed.

To solve for the optimal matching between firms and workers note that the production function $z(\tilde{s}, y)$ is assumed to be supermodular. By the well known Becker-Shapley-Shubik result (Becker, 1973 and Shapley and Shubik, 1971) the optimal matching in the labor market is positive assortative between y and $\tilde{s}$, so matching function μ is increasing in $\tilde{s}$. Moreover, the wage function w is derived from the firms' optimality condition (2), evaluated at the optimal matching μ. In the quantitative model where G and $\tilde{N}_s$ are discrete, we approximate the integral in (2) numerically, using trapezoidal integration.

The output from solving the equilibrium in the labor market given marital and household choices is the tuple (μ, w) as defined above.

B.2 Optimal Household Choices ([hh])

Second, we derive the solution of the household problem that yields spouses' optimal private consumption, (c_f, c_m), their optimal labor supply (h_f, h_m), and the distribution of effective types $\tilde{N}_s$.

Individuals arrive at the household stage either as singles with human capital s_i or in a couple with human capital bundle (s_f, s_m). We denote the household human capital type by two-dimensional vector $\mathbf{s} = (s_f, s_m) \in \{\mathcal{S} \cup \emptyset\}^2$ where, for example, $(s_f, \emptyset)$ denotes the household of single woman of type s_f.

When solving their household problem agents *take as given* wage function w, the marriage market matching distribution η, and the marriage market clearing price v.

Given prices and marriage outcomes, couples solve problem (19) and singles solve problem (20). Replacing the constraints in the objective function and noting the transferable utility structure of the problem, the collective problem of couple (s_f, s_m) after labor supply preference shocks realize is given by:

$$\max_{h_m, h_f} w(\tilde{s}_m) + w(\tilde{s}_f) + 2p^M(1 - h_m, 1 - h_f) + \delta^{h_m} + \delta^{h_f} \tag{28}$$

where $w(\tilde{s}_m)$ and $w(\tilde{s}_f)$ depend on hours through the effective human capital types (21).

Similarly, the problem of a single woman of type s_f after realization of her labor supply preference shock is

$$\max_{h_f} w(\tilde{s}_f) + p^U(1 - h_f) + \delta^{h_f} \tag{29}$$

and the problem of a single man s_m is given by

$$\max_{h_m} w(\tilde{s}_f) + p^U(1 - h_m) + \delta^{h_m}. \tag{30}$$

To derive aggregate labor supply and the distribution of effective types $\tilde{N}_s$, we need to introduce some notation.

We denote the *alternative* of hours that a decision maker chooses by $\mathbf{h} \in \{\mathcal{H} \cup \emptyset\}^2 := \{\{0, ..., 1\} \cup \emptyset\}^2$ (where $\emptyset$ indicates the hours of the non-existing partner when the individual is single). We denote by $\mathbf{h}^t$ the hours alternative chosen by a decision maker of type $t \in \{M, U\}$:

$$\mathbf{h}^t = \begin{cases} (h_i, \emptyset), i \in \{f, m\} & \text{if } t = U \\ (h_f, h_m) & \text{if } t = M. \end{cases}$$

where type $t = U$ indicates *single* (or Unmarried) and type $t = M$ indicates *couple* (or Married).

Also, we denote the economic utility associated with hours alternative $\mathbf{h}^t$ of household type $t \in \{M, U\}$ with human capital type $\mathbf{s} \in \{\mathcal{S} \cup \emptyset\}^2$ by $\overline{u}^t_{\mathbf{s}}(\mathbf{h}^t)$, where

$$\overline{u}^t_{\mathbf{s}}(\mathbf{h}^t) = \begin{cases} w(\tilde{s}_i) + p^U(1 - h_i) & \text{if } t = U \\ w(\tilde{s}_m) + w(\tilde{s}_f) + 2p^M(1 - h_m, 1 - h_f) & \text{if } t = M. \end{cases} \tag{31}$$

We obtain the optimal labor supply and private consumption (c_m, c_f, h_m, h_f) for each household by solving problems (28)-(30). Given our assumption that the labor supply shock distribution is Type-I extreme value, we then obtain the fraction of agents that optimally chooses each hours alternative. The probability that household type $t \in \{M, U\}$ with human capital type $\mathbf{s} \in \{\mathcal{S} \cup \emptyset\}^2$ chooses hours alternative $\mathbf{h} \in \{\mathcal{H} \cup \emptyset\}^2$ is

$$\pi^t_{\mathbf{s}}(\mathbf{h}^t) \quad = \quad \frac{\exp(\overline{u}^t_{\mathbf{s}}(\mathbf{h}^t)/\sigma_\delta)}{\sum_{\tilde{\mathbf{h}} \in \{\mathcal{H} \cup \emptyset\}^2} \exp(\overline{u}^t_{\mathbf{s}}(\tilde{\mathbf{h}}^t)/\sigma_\delta)} \tag{32}$$

Denoting the fraction of *households* who are type $\mathbf{s}$ by $\eta_{\mathbf{s}}$, the fraction of households who are of type $\mathbf{s}$ *and* choose hours alternative $\mathbf{h}$ is given by

$$\eta_{\mathbf{s}} \times \pi^t_{\mathbf{s}}(\mathbf{h}^t).$$

From this distribution of *household* labor supply we back out the distribution of *individual* labor supply. To do so, we compute the fraction of men and women of each *individual* human capital type, s_i in household $\mathbf{s}$, optimally choosing each individual hours alternative h_i associated with household labor supply $\mathbf{h}$. Given the distribution of individual labor supply, we can compute the distribution of effective human capital types, $\tilde{N}_s$. First, note that the support of the distribution is obtained by applying functional forms (21) for any combination of *individual* hours and skill types. Second, to each point in the support of $\tilde{s}$ we attach the corresponding frequencies from the individual labor supply distribution backed out as explained above.

Given (w, μ, η), the output from solving the household problem is the tuple $(h_f, h_m, c_f, c_m, \tilde{N}_s)$.

B.3 Partial Equilibrium in the Marriage Market ([mpe])

In the marriage stage, individuals draw idiosyncratic taste shocks for partners and single-hood, β_i^s, with $i \in \{f, m\}$ and $s \in \{\mathcal{S} \cup \emptyset\}$. At this stage, labor supply shocks are not yet realized. As a result, the *ex ante* economic value from a marriage of type (s_f, s_m) is the expected value of (28); and the ex-ante economic value from female and male single-hood is the expected value of (29) and (30). In both cases, the expectation is taken over the distribution of δ-shocks. Denoting the utility transfer to a female spouse of type s_f by $v(s_f)$, the values of being married (economic plus non-economic) for a female type s_f and a male type s_m in couple (s_f, s_m) are given by

$$\begin{aligned}
\Phi_f(s_m, s_f, v(s_f)) + \beta_f^{s_m} &:= v(s_f) + \beta_f^{s_m} \\
\Phi_m(s_m, s_f, v(s_f)) + \beta_m^{s_f} &:= \mathbb{E}_\delta \Big\{ \max_{h_m, h_f} w(\tilde{s}_m) + w(\tilde{s}_f) + 2p^M(1 - h_m, 1 - h_f) + \delta^{h_m} + \delta^{h_f} \Big\} + \beta_m^{s_f} - v(s_f) \\
&= \sigma_\delta \Big[\kappa + \log \Big(\sum_{\mathbf{h}^M \in \mathcal{H}^2} \exp\{\overline{u}_{\mathbf{s}}^M(\mathbf{h}^M)/\sigma_\delta\} \Big) \Big] + \beta_m^{s_f} - v(s_f)
\end{aligned}$$

where $\kappa = 0.57722$ is the Euler constant, $\overline{u}$ is defined in (31) and $\mathbb{E}_\delta$ indicates that the expectation is taken over the distribution of δ-shocks.

In turn, the value of being single for woman s_f is

$$\Phi_f(\emptyset, s_f) + \beta_f^{\emptyset} := \sigma_\delta \Big[\kappa + \ln \Big(\sum_{\mathbf{h}^U \in \mathcal{H}} \exp\{\overline{u}_{(s_f, \emptyset)}^U(\mathbf{h}^U)/\sigma_\delta\} \Big) \Big] + \beta_f^{\emptyset}$$

and for man s_m it is

$$\Phi_m(s_m, \emptyset) + \beta_m^{\emptyset} := \sigma_\delta \Big[\kappa + \ln \Big(\sum_{\mathbf{h}^U \in \mathcal{H}} \exp\{\overline{u}_{(\emptyset, s_m)}^U(\mathbf{h}^U)/\sigma_\delta\} \Big) \Big] + \beta_m^{\emptyset}.$$

Every man with type s_m and every woman with type s_f then chooses the skill type of their partner or to remain single to maximize their value on the marriage market:

$$\begin{aligned}
\max \quad & \{\max_{s_f \in \mathcal{S}} \Phi_m(s_m, s_f, v(s_f)) + \beta_m^{s_f}, \Phi_m(s_m, \emptyset) + \beta_m^{\emptyset}\} \\
\max \quad & \{\max_{s_m \in \mathcal{S}} \Phi_f(s_m, s_f, v(s_f)) + \beta_f^{s_m}, \Phi_f(\emptyset, s_f) + \beta_f^{\emptyset}\}
\end{aligned}$$

In practice, using the transferable utility property of our model, we solve for the *optimal* marriage matching by maximizing the total sum of marital values across all individuals in the economy, using a linear program. We denote the matching distribution by η, which solves

$$\max_{\eta_{(s,s')}\in[0,1]} \sum_{(s,s')\in\{\mathcal{S}\cup\emptyset\}^2} \eta_{(s,s')} \times \left(\Phi(s,s') + \tilde{\beta}\right)$$

$$s.t. \quad \sum_{s\in\mathcal{S}} \eta_s = 1/2$$

$$\sum_{s'\in\mathcal{S}} \eta_{s'} = 1/2$$

where $\eta_{(s,s')}$ denotes the mass of household type $(s,s') \in \{\mathcal{S}\cup\emptyset\}^2$ under matching η; η_s denotes the marginal distribution of η with respect to the first dimension; $\eta_{s'}$ denotes the marginal distribution of η with respect to the second dimension; $\Phi(s,s')$ denotes the economic value from marriage for the different types of households, $\Phi(s,s') \in \{\Phi_m(s,s',v(s')) + \Phi_f(s,s',v(s')), \Phi_f(\emptyset,s'), \Phi_m(s,\emptyset)\}$; and $\tilde{\beta}$ denotes $\beta_f^{s_m} + \beta_m^{s_f}$ for couples and $\beta_i^{\emptyset}$ $(i = \{f,m\})$ for singles. Note that the restrictions of this linear program impose that the mass of women and men across all households (couples or singles) must be equal to the total mass of women and men in the economy (which is 1/2 for both sexes).

We obtain the equilibrium matching in the marriage market, η, by solving this linear program, taking prices and allocations in households and the labor market, (w, μ, h_f, h_m), as given.

B.4 General Equilibrium of the Model

Once we have derived the solution of each of the stages taking the output from the other stages as given, we solve for the *general equilibrium* of the model by searching for the prices, allocations, and assignments such that all markets are simultaneously in (partial) equilibrium. To preview, we start "backwards" from the output of the labor market stage with an arbitrary initial wage function indicating a wage offered to each effective type. In the household stage, each *potential* household takes those wages as given and makes their labor supply choices. These optimal labor supply choices (in each potential household) are then used by each individual in the marriage market to compute the value of single-hood and marriage with different partners, leading to marriage choices. The hours choices of *formed* households give rise to a distribution of *effective* types. With this endogenous distribution of effective types we go back to the labor market stage, where we match workers' effective types with firms' productivities optimally. This labor market matching gives rise to a new wage function supporting this allocation. With the new wage function in hand, we solve and update the household and marriage problems and iterate in this manner until convergence of the wage function, i.e. until we have found the fixed point in the wage function.

B.4.1 Trembling Effective Types

A challenge in the search for the equilibrium is that each household type needs to face a wage for *any* hours choice in order to make its optimal labor supply choices. However, it may be the case that at a given iteration of our fixed point algorithm, the wage function is such that certain levels of hours are *not* chosen by some household types. Therefore, in the next iteration, agents would face a wage function that only maps *realized* effective types to a wage (i.e. a wage function 'with gaps in the support'), see subsection B.1. The problem then is that agents do not know the payoff from all potential hours choices when they try to make their optimal choice.

To fill in the gaps so that households of each type observe wages for *any* hours alternatives, we develop a trembling strategy. The trembling strategy consists of drawing a random sample of women and men and force them to supply a suboptimal amount of hours from the set of unchosen hours in each iteration. In practice, for each group of women with skill type s_f and each group of men with skill type s_m, we track their optimal choices for a given wage function and determine the hours that were *not* chosen with positive probability. We then draw a 1% random sample of women and men within each of those skill types and assign them uniformly to the unchosen hours. So we force maximally 1% of each skill type (the 'tremblers') to choose sub-optimal hours, or in other words, to tremble. Finally, we construct the distribution of effective types $\tilde{N}_s$ by taking into account both 'trembling' effective types and 'optimal' effective types.

B.4.2 Fixed Point Algorithm

To solve for the general equilibrium we denote by $\tilde{N}^*_s$ the distribution of *realized* effective types (based on *optimal* hours choices, not *trembling* hours choices). Similarly, we denote by w^* the wage as a function of *realized* effective types only, where recall that the *full support* wage function is denoted by w. The fixed point algorithm we designed to solve for the equilibrium is as follows:

0. Initiate a round-zero wage function for all possible effective types, w^0.

At any round $r \geq 1$

1. Input w^{r-1} and solve [hh] and [mpe]. Update $\tilde{N}^{*r}_s$.
2. Input $\tilde{N}^{*r}_s$ and solve [lpe]. Update w^{*r}.
3. Update w^r:
   (a) We determine w^{*r} from step 2. above.
   (b) Simultaneously, we fill in the wage for effective types that did not realize at round r by solving step 2. for *trembling types*, yielding w^r.
4. Move to round $r+1$ by going back to step 1. above and continue iterating until the wage function converges, that is, $w^{r+1}(\tilde{s}) - w^r(\tilde{s}) < \epsilon$ for $\epsilon > 0$ and small, element-by-element (for each $\tilde{s}$).
5. (OUTPUT) Compute the general equilibrium as the tuple of outputs from [hh], [mpe], and [lpe] at the round where the wage function w^r converged.

C Estimation

C.1 Identification

Identification of the Worker and Job Distribution. We identify the distributions (G, N_s) directly from the data. We treat the distribution of occupational attributes G as observable. We identify the workers' human capital distribution N_s from workers' education and fixed effect in a panel wage regression. See Section 7.3 for the details on estimation.

C.1.1 Proof of Proposition 3

Identification of the Production Function. We follow arguments on the estimation of hedonic models to show identification of the production function z. In principle, this argument is non-parametric, but in line with our parametric estimation, we focus here on the parametric approach. We mainly follow Ekeland et al. (2004), Section IV.D, and also make use of their discussion of the identification strategy proposed by Rosen (1974) and criticized by Brown and Rosen (1982). The identification is based on the firm's FOC and exploits the non-linearity of our matching model, which is an important source of identification just as in Ekeland et al. (2004). Recall the firm's optimality condition satisfies:

$$w'(\tilde{s}) = z_{\tilde{x}}(\tilde{s}, \mu(\tilde{s})) \tag{33}$$

This equation can be used to identify the parameters of interest. There are two steps:

1. Estimate the marginal return $w'(\tilde{s})$ as the derivative of the kernel regression of w (observed) on $\tilde{s}$. Denote this estimate by $\widehat{w}_{\tilde{s}}$. We treat the derivative of the wage as observable. Also note that we only observe $\tilde{s}$ for men in the data (for women, there is – at this stage – an unknown productivity wedge ψ), and so for this argument we focus on the subsample of men.

2. Estimate FOC (33) after applying a log transformation and taking into account measurement error:

$$\log(\widehat{w}_{\tilde{s}}(\tilde{s})) = \log(z_{\tilde{x}}(\tilde{s}, \mu(\tilde{s}))) + \epsilon \tag{34}$$

   where, for concreteness, we assume the functional form for z, $z(\tilde{s}, y) = A_z \tilde{s}^{\gamma_1} y^{\gamma_2} + K$ (see main text), and where we treat $\tilde{s}$ and the matching μ as observed. Note that this functional form of z circumvents the identification problem of Rosen (1974), discussed in Brown and Rosen (1982) and Ekeland et al. (2004), since the slope of the wage gradient in $\tilde{s}$ is *not* equal to the slope of the marginal product in $\tilde{s}$. We assume that ϵ is the measurement error of the marginal wage, with mean zero and uncorrelated with the right-hand-side (RHS) variables. Regression (34) identifies $(A_z, \gamma_1, \gamma_2)$.

In turn, the constant in the production function K is identified from the wage of the lowest productive type $s = 0$ (and thus $\tilde{s} = 0$), with $w(0) = \int_0^0 z_1(t, \mu(t))dt + K = K$.

Identification of the Female Productivity Wedge. We can identify ψ from the within sh-type (agents with the same s and same work hours h) wage gap across gender. Denote the gender wage gap within individuals of hours-human-capital type $sh = \hat{s}\hat{h}$ in the data by $gap(\hat{s}\hat{h})$, which we treat as observed for any $\hat{s}\hat{h}$. We here focus on any 'interior' type with $\hat{h} > 0$. Moreover, to ease exposition, we focus on identifying $\psi \in [0, 1]$, as this is the empirically relevant case (but the argument can be extended to $\psi > 1$).

Then, given the wage function and our assumption that effective skill types of women and men are given by $\tilde{s}_f = \psi s_f h_f$ and $\tilde{s}_m = s_m h_m$, the observed gender wage gap at $\hat{s}\hat{h}$ can be expressed as:

$$gap(\hat{s}\hat{h}) = \frac{w(\hat{s}\hat{h}) - w\left(\psi\hat{s}\hat{h}\right)}{w(\hat{s}\hat{h})},$$

where we made the dependence of the female wage on ψ explicit. Note that (G, N_s) were identified directly from the data and so we observe which worker matches to which firm. Thus, we consider the matching μ as known at this stage.

Then, for any observed $gap(\hat{s}\hat{h})$ with $0 \leq gap(\hat{s}\hat{h}) \leq 1 - K/w(\hat{s}\hat{h})$, the female wage is given by:

$$w\left(\psi\hat{s}\hat{h}\right) = w(\hat{s}\hat{h})(1 - gap(\hat{s}\hat{h})) \tag{35}$$

For a given (observed) μ, the RHS is independent of ψ, positive and finite. In turn, the LHS is positive and finite; and it is a continuous and strictly increasing function in ψ with $w\left(\psi\hat{s}\hat{h}\right) = K$ for $\psi = 0$ and $w\left(\psi\hat{s}\hat{h}\right) = w(\hat{s}\hat{h})$ for $\psi = 1$.

Hence, one of the following is true: either there is an interior gap, $0 < gap(\hat{s}\hat{h}) < 1 - K/w(\hat{s}\hat{h})$, and so by the Intermediate Value Theorem there exists a unique $\psi \in (0, 1)$ for which (35) holds; or, a minimal gap $gap(\hat{s}\hat{h}) = 0$ pins down $\psi = 1$ or a maximal gap $gap(\hat{s}\hat{h}) = 1 - K/w(\hat{s}\hat{h})$ pins down $\psi = 0$. Thus, ψ is identified from gender wage gaps of agents with the same hours-human-capital combination.

Identification of the Scale of the Labor Supply Shock. Recall that the choice set of singles differs from that of couples. In Appendix B, we introduced the notation where we denote the *alternative* of hours that a decision maker $t \in \{M, U\}$ chooses by $\mathbf{h}^t \in \{\mathcal{H} \cup \emptyset\}^2 := \{\{0, ..., 1\} \cup \emptyset\}^2$ with:

$$\mathbf{h}^t = \begin{cases} (h_i, \emptyset), i \in \{f, m\} & \text{if } t = U \\ (h_f, h_m) & \text{if } t = M. \end{cases}$$

where type $t = U$ indicates *unmarried* and type $t = M$ indicates *married.*

Also, we denote the sum of economic utility and utility derived from preference shocks of decision-maker t with human capital type $\mathbf{s} \in \{\mathcal{S} \cup \emptyset\}^2$ by $\overline{u}^t_{\mathbf{s}}(\mathbf{h}^t) + \delta^{\mathbf{h}^t}$, where

$$\overline{u}^t_{\mathbf{s}}(\mathbf{h}^t) + \delta^{\mathbf{h}^t} = \begin{cases} u(c_i, p^U(1 - h_i)) + \delta^{h_i}, i \in \{f, m\} & \text{if } t = U \\ u(c_f, p^M(1 - h_m, 1 - h_f)) + u(c_m, p^M(1 - h_m, 1 - h_f)) + \delta^{h_f} + \delta^{h_m} & \text{if } t = M. \end{cases}$$

The probability that household type t with human capital $\mathbf{s}$ chooses hours alternative $\mathbf{h}$ is

$$\pi^t_{\mathbf{s}}(\mathbf{h}^t) \quad = \quad \frac{\exp(\overline{u}^t_{\mathbf{s}}(\mathbf{h}^t)/\sigma_\delta)}{\sum_{\tilde{\mathbf{h}}\in\{\mathcal{H}\cup\emptyset\}^2}\exp(\overline{u}^t_{\mathbf{s}}(\tilde{\mathbf{h}}^t)/\sigma_\delta)} \tag{36}$$

which follows from our assumption on the preference shock distribution (Type-I extreme value).

Let $\mathbf{h}^U = \mathbf{0} := (0, \emptyset)$ denote the hours for a single who puts all available time into home production and works zero hours in the labor market. We consider alternative $\mathbf{h}^U = \mathbf{0}$ as our normalization choice and obtain for a single male of human capital type $\mathbf{s} = (s_m, \emptyset))$ the relative choice probabilities:

$$\begin{aligned}
\frac{\pi^U_{\mathbf{s}}(\mathbf{h}^U)}{\pi^U_{\mathbf{s}}(\mathbf{0})} &= \frac{\exp(\overline{u}^U_{\mathbf{s}}(\mathbf{h}^U)/\sigma_\delta)}{\exp(\overline{u}^U_{\mathbf{s}}(\mathbf{0})/\sigma_\delta)} \\
\log\left(\frac{\pi^U_{\mathbf{s}}(\mathbf{h}^U)}{\pi^U_{\mathbf{s}}(\mathbf{0})}\right) &= \frac{\overline{u}^U_{\mathbf{s}}(\mathbf{h}^U) - \overline{u}^U_{\mathbf{s}}(\mathbf{0})}{\sigma_\delta} \\
&= \frac{w(s_m h_m) - w(s_m 0) + p^U(1-h_m) - p^U(1-0)}{\sigma_\delta} \\
&= \frac{w(s_m h_m) + p^U(1-h_m) - p^U(1)}{\sigma_\delta} \qquad (37)
\end{aligned}$$

where the wage from not working is set to zero and where h_m is the male hours associated to this single household's hours choice, $\mathbf{h}^U = (h_m, \emptyset)$. We treat human capital types as observed at this stage and consider two single types $\mathbf{s}' = (s'_m, \emptyset)$ and $\mathbf{s}'' = (s''_m, \emptyset)$. Then we can consider the difference in relative choices of these two single men:

$$\log\left(\frac{\pi^U_{\mathbf{s}'}(\mathbf{h}^U)}{\pi^U_{\mathbf{s}'}(\mathbf{0})}\right) - \log\left(\frac{\pi^U_{\mathbf{s}''}(\mathbf{h}^U)}{\pi^U_{\mathbf{s}''}(\mathbf{0})}\right) \quad = \quad \frac{1}{\sigma_\delta}\left(w(s'_m h_m) - w(s''_m h_m)\right).$$

The LHS is observed in the data (how does the relative choice probability for hours alternative $\mathbf{h}^U \neq \mathbf{0}$ change in the population of male singles as one varies human capital s_m), and on the RHS, the wage difference (it is the effect of men's human capital on wages given the hours choice $\mathbf{h}^U \neq \mathbf{0}$) is also observed and different from zero as the wage strictly increases in human capital. Thus, σ_δ is identified.

Identification of the Home Production Function. Let $\mathbf{h}^M = \mathbf{1} := (1, 1)$ denote the vector of hours for couples in which both spouses put zero hours into home production and thus work full time in the labor market. Alternative $\mathbf{h}^M = \mathbf{1}$ is our normalization choice and we obtain the relative choice probabilities of married couple $\mathbf{s}$ of choosing hours $\mathbf{h}^M \neq \mathbf{1}$ versus $\mathbf{h}^M = \mathbf{1}$ as:

$$\begin{aligned}
\frac{\pi^M_{\mathbf{s}}(\mathbf{h}^M)}{\pi^M_{\mathbf{s}}(\mathbf{1})} &= \frac{\exp(\overline{u}^M_{\mathbf{s}}(\mathbf{h}^M)/\sigma_\delta)}{\exp(\overline{u}^M_{\mathbf{s}}(\mathbf{1})/\sigma_\delta)} \\
\log\left(\frac{\pi^M_{\mathbf{s}}(\mathbf{h}^M)}{\pi^M_{\mathbf{s}}(\mathbf{1})}\right) &= \frac{\overline{u}^M_{\mathbf{s}}(\mathbf{h}^M) - \overline{u}^M_{\mathbf{s}}(\mathbf{1})}{\sigma_\delta} \\
&= \frac{w(\psi s_f h_f) - w(\psi s_f) + w(s_m h_m) - w(s_m) + 2p^M(1-h_m, 1-h_f)}{\sigma_\delta} \qquad (38)
\end{aligned}$$

where we used that $2p^M(0,0) = 0$ by assumption in our quantitative model. Note that the LHS of (38) (relative choice probabilities) is observed, and on the RHS, wages of men and women with types (s_f, s_m) conditional on hours are also observed in the data, and σ_δ is known at this stage. Thus, home production function p^M is non-parametrically identified since we can specify (38) for all hours alternatives $\mathbf{h}^M$ chosen in the data. Note that we can identify p^M from a couple of any type $\mathbf{s} = (s_f, s_m)$.

By a similar argument the home production function of singles, p^U, is identified.

Identification of the Scale of the Marriage Taste Shock. In this section, we show that σ_β^M is identified once the parameters of the utilities are identified, where we will impose the following (mild) assumption.

Assumption D1 (Identification)**.**
(Marriage Sorting.) There exists male types $(s'_m, s''_m) \in \mathcal{S}^2$*,* $s''_m > s'_m$*, and female types* $(s'_f, s''_f) \in \mathcal{S}^2$*,* $s''_f > s'_f$*, such that* $\frac{\eta_{(s''_f, s''_m)}}{\eta_{(s'_f, s''_m)}} \neq \frac{\eta_{(s''_f, s'_m)}}{\eta_{(s'_f, s'_m)}}$.

The assumption states that there is marriage market sorting, at least somewhere in the support of (s_f, s_m). To see this, consider the following case for illustration: Suppose that for all $s''_m > s'_m$, and $s''_f > s'_f$, $\frac{\eta_{(s''_f, s''_m)}}{\eta_{(s'_f, \hat{s}''_m)}} > \frac{\eta_{(s''_f, s'_m)}}{\eta_{(s'_f, \hat{s}'_m)}}$. In words, $\eta_{(s_f, s_m)}$ has the monotone likelihood ratio property, or equivalently, is log-supermodular. Then, higher male types s_m are matched with higher female types s_f in the first-order stochastic dominance sense, i.e. there is positive sorting in s-types.

Let $\eta_{(s_f, s_m)}$ be the probability that a man s_m chooses woman s_f on the marriage market, conditional on marrying. Under the assumption that the taste shock is extreme-value distributed (and following the same derivations as for the choice probabilities of hours), $\eta_{(s_f, s_m)}$ is given by:

$$\eta_{(s_f, s_m)} = \frac{\exp(\Phi(s_m, s_f, v(s_f)))/\sigma_\beta^M)}{\sum_{s'_f} \exp(\Phi(s_m, s'_f, v(s'_f)))/\sigma_\beta^M)}$$

where, as before, we denote by $\Phi(s_m, s_f, v(s_f))$ the expected value of man s_m from being married to woman s_f and paying her the transfer $v(s_f)$. This value is given by:

$$\begin{aligned}
\Phi(s_m, s_f, v(s_f)) &:= \sigma_\delta \Big[\kappa + \log\Big(\sum_{\mathbf{h}^M \in \mathcal{H}^2} \exp\Big\{\overline{u}_{\mathbf{s}}^M(\mathbf{h}^M)/\sigma_\delta\Big\}\Big)\Big] - v(s_f) \\
&= \sigma_\delta \Big[\kappa + \log\Big(\sum_{\mathbf{h}^M \in \mathcal{H}^2} \exp\Big\{\frac{w(s_m h_m) + w(\psi s_f h_f) + 2p^M(1-h_m, 1-h_f)}{\sigma_\delta}\Big\}\Big)\Big] - v(s_f)
\end{aligned}$$

Using the ratio of probabilities of choosing two different women s''_f and s'_f, we obtain:

$$\log\left(\frac{\eta_{(s''_f, s_m)}}{\eta_{(s'_f, s_m)}}\right) = \frac{\Phi(s''_f, s_m, v(s''_f)) - \Phi(s'_f, s_m, v(s'_f))}{\sigma_\beta^M}.$$

It follows from Assumption D1, that there exist types $(s'_f, s''_f, s'_m, s''_m)$ such that

$$\log\left(\frac{\eta_{(s''_f,s''_m)}}{\eta_{(s'_f,s''_m)}}\right) - \log\left(\frac{\eta_{(s''_f,s'_m)}}{\eta_{(s'_f,s'_m)}}\right) = \frac{\Phi(s''_f, s''_m, v(s''_f)) - \Phi(s'_f, s''_m, v(s'_f))}{\sigma_\beta^M} - \frac{\Phi(s''_f, s'_m, v(s''_f)) - \Phi(s'_f, s'_m, v(s'_f))}{\sigma_\beta^M} \neq 0.$$

which, using the expression for $\Phi(s_m, s_f, v(s_f))$ from above, we can spell out as:

$$\begin{aligned}
\log\left(\frac{\eta_{(s''_f,s''_m)}}{\eta_{(s'_f,s''_m)}}\right) - \log\left(\frac{\eta_{(s''_f,s'_m)}}{\eta_{(s'_f,s'_m)}}\right) = \frac{\sigma_\delta}{\sigma_\beta^M}\Bigg(& \log\Big(\sum_{\mathbf{h}^M\in\mathcal{H}^2}\exp\Big\{\frac{w(s''_m h_m) + w\left(\psi s''_f h_f\right) + 2p^M(1-h_m, 1-h_f)}{\sigma_\delta}\Big\}\Big) \\
& - \log\Big(\sum_{\mathbf{h}^M\in\mathcal{H}^2}\exp\Big\{\frac{w(s''_m h_m) + w\left(\psi s'_f h_f\right) + 2p^M(1-h_m, 1-h_f)}{\sigma_\delta}\Big\}\Big) \\
& - \log\Big(\sum_{\mathbf{h}^M\in\mathcal{H}^2}\exp\Big\{\frac{w(s'_m h_m) + w\left(\psi s''_f h_f\right) + 2p^M(1-h_m, 1-h_f)}{\sigma_\delta}\Big\}\Big) \\
& + \log\Big(\sum_{\mathbf{h}^M\in\mathcal{H}^2}\exp\Big\{\frac{w(s'_m h_m) + w\left(\psi s'_f h_f\right) + 2p^M(1-h_m, 1-h_f)}{\sigma_\delta}\Big\}\Big)\Bigg) \\
\neq 0 &
\end{aligned}$$

Since the LHS is not zero by assumption, the RHS is also not zero. Moreover, all objects on the RHS are either observed (wages) or identified at this stage (home production function and σ_δ, except σ_β^M). We can solve this equation for σ_β^M, giving a unique solution. Thus, σ_β^M is identified.

C.2 Estimation of Worker Types

Sample Selection. Our sample consists of individuals in the GSOEP from 1984-2018 who are between 20 and 60 years old and are either married/cohabiting or single. We exclude individual-year observations when the individual indicated self-employment and when they worked in poorly defined occupations (`kldb92`$\geq$ 9711). We also exclude observations with missing information on education or with missing (not zero) labor force experience. Our panel consists of around 212,000 person-year observations.

Key Variables. For *weekly hours*, we use reported actual hours which, when positive, we winsorized by 10 hours from below and 60 hours from above to deal with outliers. For *labor force experience* we use the reported labor force experience, and we impute it by potential experience if this information is missing.[31] For *education*, we use three categories: In the group of *low education*, there are those whose highest degree is lower secondary, high school or vocational with weakly less than 11 years of education (around 35%). In the group of *medium education*, there are those with vocational degree and above 11 years of education (around 44%). In the group of *high education*, there are those with college degree

[31] For men, potential and actual experience are almost perfectly correlated, which is why this imputation should work well for them. For women, the correlation is much lower, which is why we do not impute here.

or more (around 20%). For *occupation codes*, we use the variable `kldb92_current`, which consistently codes occupations across the entire panel. Our wage variable are *log hourly wages*, inflation-adjusted in terms of 2016 Euros. For the definition of 'demographical cells' in the selection stage below, we additionally use a variable that indicates whether *children below 3 years old* are in the household, *age bins* (≤ 25, > 25 and ≤ 40, > 40 and ≤ 50, > 50) and the *state* of residence.

Selection Equation. To account for selection into labor force participation in the wage regression, we first run a selection regression. To do so, we need an instrument that affects participation but is excluded from wage regression (22). Since the variation in participation in our sample is mainly due to women, we use the 'progressiveness' in an individual's narrowly defined demographic cell. We proxy by the share of females working in a narrowly defined demographic cell. Our cells are defined by a combination of state, year, age and an indicator whether a child under the age of 3 is in the household.[32] When defining this variable for a particular individual, we employ the 'leave-one-out' method and do not count the individual's labor force participation when computing this statistic. We further drop cells with less than five observations. We end up with around 2,500 cells with more than 5 observations. Note that we experimented with additional cell characteristics (education and country of origin) but there is a trade-off between number of observations by cell and making the cells more specific to the demographic groups. Defining the cells by these additional variables would imply to drop more than twice as many observations due to small cell size.

Our assumption on this IV is that the following exclusion restriction holds: 'progressiveness of an individual's demographic cell' only affects wages through labor force participation but not in other ways.

We run the following probit selection regression:

$$emp_{it} = \alpha\, share_{j(i)t} + \sum_{ed \in \{voc,c\}} \alpha^{ed} x_{it}^{ed} + \beta_z' Z_{it} + \kappa_s + \rho_t + \epsilon_{it} \tag{39}$$

where the dependent variable is an indicator of whether individual i is employed at time t, $share_{j(i)t}$ is the progressiveness measure in the demographic cell j of individual i at time t (given by the share of women working in the cell, see description of this IV above), x_{it}^{ed} captures education indicators (indicator variables for medium and high education, so that low education is the reference group), Z_{it} is a vector of demographic individual controls (linear and quadratic labor force experience in years, household size) and κ_s and ρ_t are state and year fixed effects and ϵ_{it} is a mean-zero error term. We cluster standard errors on the cell level.

The results are in Table 5, where we label our IV $share_{j(i)t}$ by *Share of Working Women in Cell*. There is a strong positive effect of the share of women working in the demographic cell on labor force participation of an individual in that cell (coefficient of 1.355 with standard error 0.0284).

[32] German children start kinder-garden when they are 3 years old. Before age 3, children are predominantly at home (in 2013, only 29% of children aged 0-2 were in daycare, Source: OECD (2016)), so age 3 is an important threshold when it comes to mothers' labor force participation.

Table 5: Selection Regression

	(1) Employed
Share of Working Women in Cell	1.355***
	(0.0284)
Experience	0.0919***
	(0.00216)
Experience^2	-0.00164***
	(0.0000599)
Medium Educ	0.453***
	(0.00932)
High Educ	0.851***
	(0.0115)
HH Size	-0.0462***
	(0.00539)
Constant	-0.918***
	(0.0358)
Observations	212,894
Pseudo R^2	0.144

Notes: A *Cell* is defined by state, age, year and presence of children under 3 in the household. The regression includes year and state fixed effects. Standard errors clustered at cell level in parentheses. ***Significant at the 1% level.

Wage Regression. We further restrict the sample to those that are employed and have non-missing hourly wage. Since we instrument hours worked and hours worked squared by (i) the hours worked by the partner, (ii) the hours worked by the partner squared and (iii) whether partner is present (where we set partner's hours to zero in both cases, if partner is present but not working and if partner is not present). We therefore drop observations whose partner reports to be employed but has zero reported labor hours and observations whose partner has missing employment and hours information. Since we include individual fixed effects we also drop singleton observations (those who only show up in a single year of the panel). This leads to a sample of 133,214 person-year observations. Based on our model wage function, we choose the following regression specification:

$$\ln w_{it} = \nu_i + \beta_0 IMR_{it} + \beta_1 h_{it} + \beta_2 h_{it}^2 + \sum_{ed \in \{voc,c\}} \alpha^{ed} x_{it}^{ed} + \beta_z' Z_{it} + \kappa_s + \rho_t + \epsilon_{it} \tag{40}$$

where ν_i is a person-fixed effect, IMR_{it} is the inverse mills ratio of individual i in year t from the selection probit regression, h_{it} are weekly hours worked, and x_{it}^{ed}, Z_{it}, κ_s and ρ_t are as in the selection regression (39). Note that we could not include h_{it} into the selection equation since $h_{it} = 0$ versus $h_{it} > 0$ is a perfect predictor of employment. Nevertheless, based on our model, it is important to control for hours worked in the hourly wage regression. We again cluster standard errors on the cell level.

Table 6 contains the results. In column (1) an (2) we report the first stage regressions (for two variables to be instrumented: weekly hours and weekly hours squared) and column (3) contains the second stage regression. The three IV's for the hours worked variables (partner's hours, partner's hours squared and partner present) are not subject to the weak instrument problem according to the F-statistics. Regarding the second stage, we note that the inverse mills ratio is positive and significant, indicating that individuals are positively selected into working and not controlling for selection here would have biased the coefficients upward. Moreover, we note that weekly hours worked have a strong positive effect on wages, justifying our model assumption that hours affect productivity and thus hourly wages. In particular, increasing hours from 30 to 40 hours per week yields an hourly wage return of $(40 \times 0.119 - 40^2 \times 0.00164) - (30 \times 0.119 - 30^2 \times 0.00164) = 0.042$, so of around 4%.

Table 6: Wage Regression

	(1) Weekly Hours Worked	(2) Weekly Hours Worked^2	(3) Log Hourly Wage
Partner's Weekly Hours Worked	-0.0574***	-4.607***	
	(0.00480)	(0.389)	
Partner's Weekly Hours Worked^2	0.000895***	0.0796***	
	(0.000105)	(0.00857)	
Partner Present	1.133***	60.85***	
	(0.164)	(12.57)	
Experience	0.377***	28.53***	0.0332***
	(0.0251)	(1.718)	(0.00182)
Experience^2	-0.00460***	-0.362***	-0.000555***
	(0.000399)	(0.0298)	(0.0000286)
Medium Educ	0.242	27.28*	0.0264**
	(0.216)	(15.89)	(0.0103)
High Educ	4.289***	299.8***	0.195***
	(0.458)	(30.93)	(0.0244)
HH Size	-1.280***	-80.32***	0.0218***
	(0.0510)	(3.495)	(0.00467)
Inverse Mills Ratio	-3.519***	-215.2***	0.133***
	(0.255)	(18.38)	(0.0197)
Weekly Hours Worked			0.119***
			(0.0156)
Weekly Hours Worked^2			-0.00164***
			(0.000211)
Observations	133,214	133,214	133,214
F	69.92	59.52	162.902
R^2			-0.386

Notes: All specifications include year and state fixed effects. The table reports the IV version of regression (40), where we instrument Weekly Hours Worked and Weekly Hours Worked squared by the partner's Weekly Hours Worked (linear and squared) and whether the partner is present. Columns (1) and (2) are the first stage regressions while Column (3) is the second stage regression. Standard errors (clustered at the cell level) are in parentheses. A *Cell* is defined by state, year, age, and presence of children under 3 in the household. ***Significant at the 1% level. **Significant at the 5% level.

Imputation. Based on these results, we are able to obtain x-types and ν-types (and thus s-types) for around 17,000 individuals. We impute fixed effects of the remaining ones (around 11,600 individuals) based on the *multiple imputation* approach. As auxiliary variables in this imputation we choose covariate in our data set that are most correlated with the individual fixed effects (such as *education*, *gender* and *full time labor force participation*). After imputation, we use the subset of individuals for structural estimation the comply with our final sample restrictions, Appendix OC.2. In our final estimation sample (for baseline period, 2010-2016), we have 3,857 unique individuals. For 24% of them we have imputed ν_i.

We then divide individuals into the three education groups and assess within each group whether an individual has a low (below median) or high (above median) fixed effect, so there are two subgroups in each education bin. We compute the subgroup fixed effect, $\overline{\nu}_j^{ed}$, as the mean of the individuals fixed effects belonging to subgroup j of education bin ed. This way we obtain six fixed effects groups (two for each education group). Finally, we compute the human capital type for each individual as $s_i = \alpha^{ed} x_i^{ed} + \overline{\nu}_j^{ed}(i)$ where $\overline{\nu}_j^{ed}(i)$ is the fixed effect of individual i's group. We obtain six s-types. The resulting human capital distribution (s-types by education group) is displayed in Table 7.

Table 7: Worker Distribution of s-Types by Education

	s-type (ed, $\overline{\nu}_{ed}$)						
Education	(low,low)	(med,low)	(low,high)	(med,high)	(high,low)	(high,high)	Total
Low Educ	523	0	601	0	0	0	1,124
Medium Educ	0	617	0	1,046	0	0	1,663
High Educ	0	0	0	0	447	623	1,070
Total	523	617	601	1,046	447	623	3,857

Notes: *Low Educ* includes high school and vocational education with less than 11 years of schooling. *Medium Educ* is defined as vocational education with more than 11 years of schooling. *High Educ* is defined as college and more. We obtain six s-types by ordering their value of $s_i = \alpha^{ed} x_i^{ed} + \overline{\nu}_j^{ed}(i)$ where parameters α^{ed} are estimated based on regression (40) and where $\overline{\nu}_j^{ed}(i)$ is the fixed effect of individual i in education group ed and subgroup j, also computed from model (40) as the average individual fixed effects in that group.

C.3 Estimation of Job Types

Data and Sample Selection. Our main data source for measuring occupation types is the BIBB collected in 2012 by the German Federal Institute of Vocational Training, and the German Federal Institute for Occupational Safety and Health. This survey is representative of the German employed population. In particular, it contains data on task usage in 1,235 occupations defined by the 4-digit code `kldb92`, which we also use for our analysis in the GSOEP. This data is reported by individuals who work in these occupations. In order to reduce the problem of noisy reporting, we drop occupations in which the task information is based on less than 5 individuals. We are left with task data for 613 occupations. These are the most common occupations and we will base our structural estimation exercise on the subset (608 occupations) that can be merged to the occupations held by individuals in our GSOEP sample, using the four-digit occupational codes `kldb92` from the German Classification of Occupations 1992.

Task Data. The BIBB contains data on how intensely different 4-digit occupations use different types of tasks.[33] These intensity measures are continuous and we normalize them to be on the unit interval. The reported tasks are comprised of: Detailed Work, Same Cycle, New Tasks, Improve Process,

[33] For confidentiality reasons, the BIBB Public Use Files do not contain information on the 4-digit kldb 1992 classification of occupations of each individual (it contains 3-digit levels). However, we were able to obtain summary measures of occupational task content at the 4-digit level, kldb 1992, without reference to individual identifiers.

Produce Items, Tasks not Learned, Simultaneous Tasks, Consequence of Mistakes, Reach Limits, Work Quickly, Problem Solving, Difficult Decisions, Close Gaps of Knowledge, Responsibility for Others, Negotiate, Communicate.

Model Selection Stage. We merge the task data from the BIBB into occupations held by individuals in the GSOEP. As with the worker types, we here use the entire SOEP panel (here: pooled). We run a Lasso regression of log hourly wages on the task descriptors listed in the last paragraph in order to systematically select the tasks that matter for pay. This procedure selects 13 tasks (all tasks from the dataset except: Improve Process, Consequence of Mistakes and Difficult Decisions).

Principal Component Analysis. Because we want to reduce the occupational type to a single dimension, we collapse the information of the 13 selected tasks into a single measure using a standard dimension reduction technique (principal component analysis). We then select the first principal component, which captures the most variation of the underlying task variables in the sample of employed workers in the SOEP: It captures around 43% of the underlying variation and – based on the loadings on the underlying task descriptors – our interpretation of this component is *task complexity* or *high skill requirement*. This interpretation is based on positive loadings on all task variables except *detailed work* and *same cycle*, which arguably are the only tasks in the dataset that indicate routine work. We report a scree-plot with eigenvalues of the different principal components and a plot with loadings of the first PC in Figure 10, left and right panel, respectively.

Figure 10: Principal Component Analysis

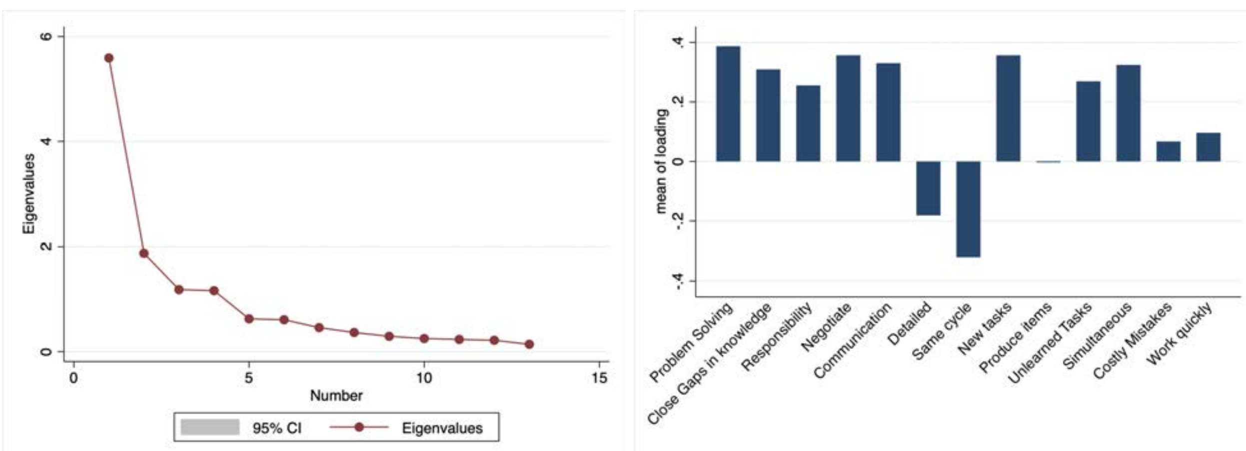

We then compute the mean of this measure by occupation and denote it by $\hat{y}$. Once matched to our main sample (see Appendix OC.2), we define our final measure of occupational type as the ranking of occupations in the task complexity distribution, i.e. $y = \widehat{G}(\hat{y})$, where $\widehat{G}$ is the cdf of $\hat{y}$. So $y \sim G$ where $G = U[0, 1]$. The reason for this transformation is that the occupational task data only has an ordinal interpretation. Note that our production function is flexible enough to capture the true output as a function of non-transformed types $\widehat{y}$. Examples of occupations in the top 5% of the G distribution include engineers and programmers. Examples of occupations in the lowest 5% include janitors and cleaners.

Table 8: Correlation Across Alternative Measures

	y (baseline)	y (PCA)	y (Lasso)
y (baseline)	1		
y (PCA)	0.983^{***}	1	
y (Lasso)	0.946^{***}	0.922^{***}	1
Observations		608	

Notes: y(baseline) is our main measure, based on two steps: Lasso and PCA.

***Significant at the 1% level.

Table 9: Summary Statistics of Alternative Measures

	mean	sd	min	max
y (baseline)	.573	.193	0	1
y (PCA)	.585	.188	0	1
y (Lasso)	.565	.196	0	1

Notes: y(baseline) is our main measure, based on two steps: Lasso and PCA.

Alternative Approaches. We considered 2 alternative approaches for the estimation of occupational types. The first alternative is to not use wage data at all but instead, to directly reduce the multi-dimensional task data to a single measure by PCA. The second alternative is to rely more heavily on wages by first selecting important tasks via Lasso and then using the *predicted* wage based on these important tasks as our measure for occupation types. With all three approaches, we get very similar results. We show the summary statistics and correlations of all three measures in Tables 9 and 8.

We also considered determining the occupational types through an occupational fixed effect in the wage regression we used to recover worker types (equation (22)). This would have meant to run a two-way fixed effects regression. We chose our alternative approach that does not rely on occupational fixed effects for the following reasons: First, based on our model featuring a competitive labor market, the wage function does not depend on an occupational fixed effect/type when controlling for workers' effective types: all workers with the same effective type $\tilde{s}$ should be matched to the same occupation. Second, the two-way fixed effects approach is known to be problematic under limited worker mobility (limited mobility bias) and when one is interested in sorting (since the correlation between worker and firm/occupation fixed effects is not accurately capturing sorting).

C.4 Internal Estimation

C.4.1 Parameters Set Outside the Model

Table 10: Exogenously Fixed Parameters

Parameter		**Value**
Hourly Minimum Wage	K	6.32
Labor Supply Shock (location)	$\bar{\delta}$	0.00
Preference Shock for Partners (location)	$\bar{\beta}^M$	0.00
Preference Shock for being Single (location)	$\bar{\beta}^U$	0.00
Preference Shock for being Single (scale)	σ_β^U	29.25

C.4.2 Results

Table 11: Targeted Moments

	Model	Data
M1. LFP Female to Male Ratio	0.7426	0.7864
M2. Full Time Work Female to Male Ratio	0.4046	0.3834
M3. LFP Married to Single Ratio, Men	0.8612	0.8556
M4. LFP Married to Single Ratio, Women	0.9890	1.2534
M5. Correlation of Spouses Home Hours	0.3159	0.3120
M6. Mean Hourly Wage	17.7271	17.6354
M7. Variance Hourly Wage	51.1067	53.9061
M8. Upper Tail (90-50) Wage Inequality	3.0852	2.9686
M9. Overall (90-10) Wage Inequality	1.7294	1.7271
M10. Correlation between Spouses Types	0.4403	0.4468
M11. Fraction of Single Men	0.2055	0.1976
M12. Gender Wage Gap by Effective Type 2	0.1227	0.1557
M13. Gender Wage Gap by Effective Type 4	0.1414	0.1464
M14. Female LFP by Couple Types 3 and 4	0.7242	0.7308
M15. Female LFP by Couple Types 5 and 6	0.8468	0.8071
M16. Female LFP of Single Women Type 3 and 4	0.8076	0.7320
M17. Female LFP of Single Women Type 5 and 6	0.8583	0.8429

Notes: *LFP* stands for Labor Force Participation. Moments are computed as discussed in Table O4.

Table 12: Un-targeted Moments: Marriage Matching Frequencies - Model and (Data)

	Low Educ Men	**Medium Educ Men**	**High Educ Men**	**Single Women**
Low Educ Women	0.0820 (0.0747)	0.0721 (0.0449)	0.0230 (0.0126)	0.0361 (0.0365)
Medium Educ Women	0.1016 (0.0860)	0.1246 (0.2159)	0.0984 (0.0695)	0.0787 (0.0747)
High Educ Women	0.0361 (0.0149)	0.0459 (0.0485)	0.0754 (0.0986)	0.0557 (0.0562)
Single Men	0.0557 (0.0527)	0.0656 (0.0714)	0.0492 (0.0430)	0.0000

Notes: *Low Educ* includes high school and vocational education with less than 11 years of schooling. *Medium Educ* is defined as vocational education with more than 11 years of schooling. *High Educ* is defined as college and more. Data frequencies are shown in parenthesis.

D Quantitative Analysis

D.1 Comparative Statics

Figure 11: a. Gender Wage Gap, b. Household Income Variance Decomposition, c. Household Income Variance, d. Marriage Market Sorting, e. Gender Gap Labor Hours, f. Gender Gap Labor Sorting.

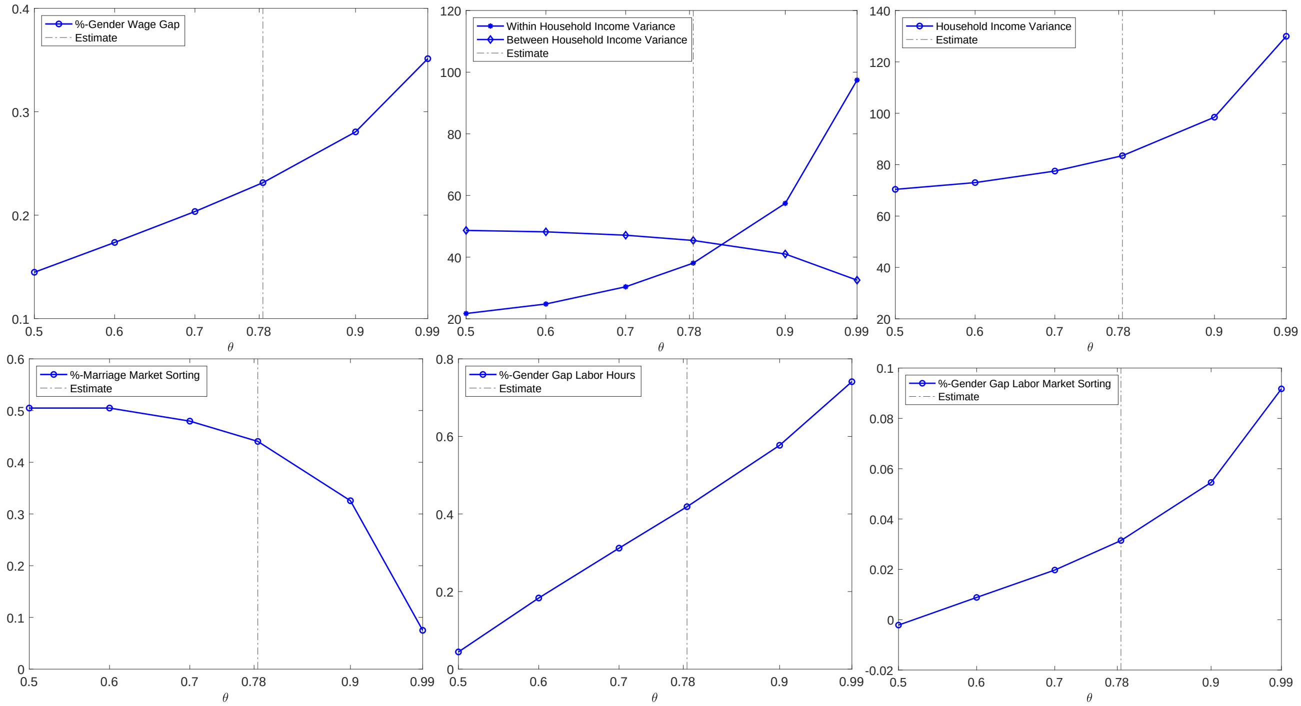

Figure 12: a. Gender Wage Gap, b. Household Income Variance Decomposition, c. Household Income Variance, d. Marriage Market Sorting, e. Gender Gap Labor Hours, f. Gender Gap Labor Sorting.

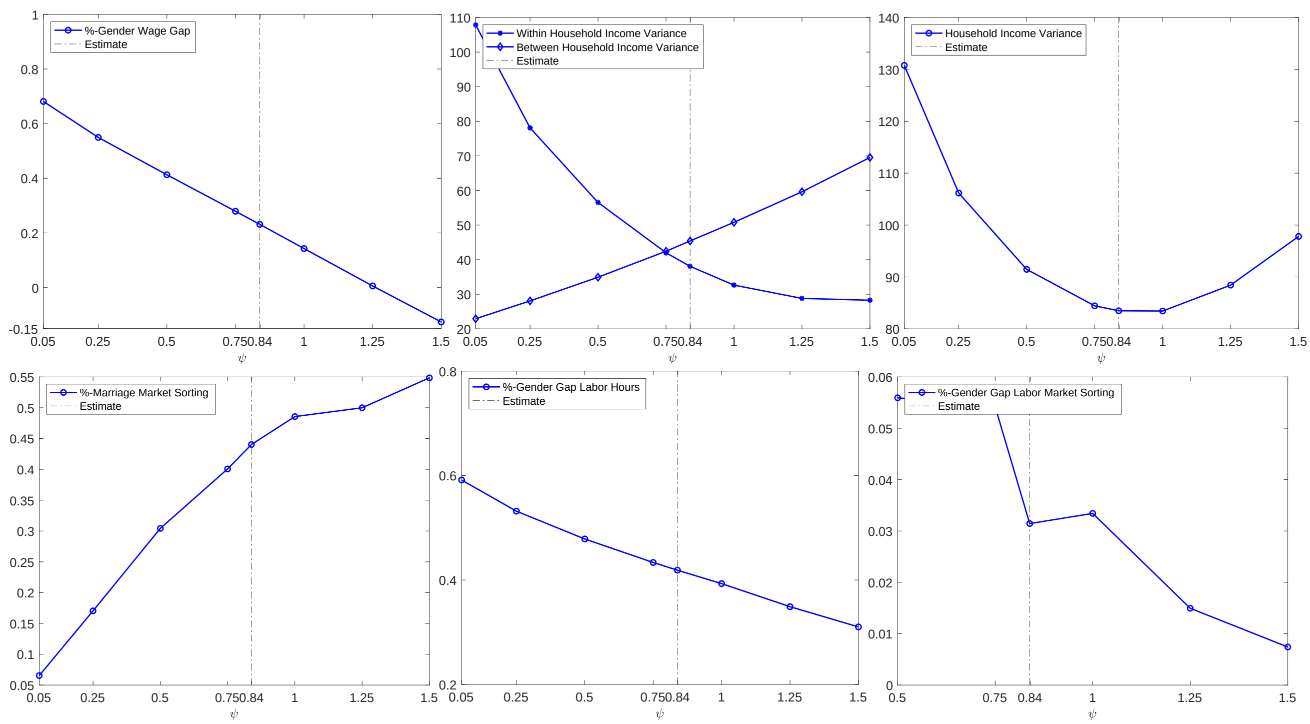

D.2 Inequality Over Time

Table 13: Data and Model moments: 1990-1996 versus 2010-2016

	Past Period		**Current Period**		**Data Diff.**
	Model	Data	Model	Data	p-value
M1. LFP Female to Male Ratio	0.6266	0.5875	0.7426	0.7864	0.0000
M2. Full Time Work Female to Male Ratio	0.2374	0.3336	0.4046	0.3834	0.2324
M3. LFP Married to Single Ratio, Men	0.7768	0.6343	0.8612	0.8556	0.0190
M4. LFP Married to Single Ratio, Women	1.0057	1.0844	0.9890	1.2534	0.0000
M5. Correlation of Spouses Home Hours	0.1672	0.1518	0.3159	0.3120	0.0019
M6. Mean Hourly Wage	16.9807	17.0106	17.7271	17.6354	0.0000
M7. Variance Hourly Wage	35.3776	37.2476	51.1067	53.9061	0.0026
M8. Upper Tail (90-50) Wage Inequality	2.5262	2.3486	3.0852	2.9686	0.0000
M9. Overall (90-10) Wage Inequality	1.5604	1.5841	1.7294	1.7271	0.0000
M10. Correlation between Spouses Types	0.3864	0.4052	0.4403	0.4468	0.0000
M11. Fraction of Single Men	0.1186	0.1147	0.2055	0.1976	0.2069
M12. Gender Wage Gap by Effective Type 2	0.1682	0.1814	0.1227	0.1557	0.0000
M13. Gender Wage Gap by Effective Type 4	0.1719	0.1839	0.1414	0.1464	0.4885

Notes: *LFP* stands for Labor Force Participation. Moments are computed as discussed above in Appendix C.4. The last column of the table reports the p-value of the hypothesis test of the differences between the data moments in the two samples being zero. We use a standard T-test for differences in means (M6), a standard Levene test for differences in variances (M7) and standard tests for differences in proportions (M11). We use a Fisher transformation to construct the test statistic for the differences in correlations between samples (M5 and M10). We use a two-sample Wald test for differences in ratios across samples (M1, M2, M3, M4, M8, M9, M12 and M13). To construct the statistic for the Wald tests for difference in ratios, we use bootstrap techniques for the variance estimation.

Table 14: Estimated Parameters: 1990-1996 versus 2010-2016

		Past Period		**Current Period**	
		Estimate	s.e.	Estimate	s.e.
Female Relative Productivity in Home Production	θ	0.88	0.01	0.78	0.02
Complementarity Parameter in Home Production	ρ	-0.16	0.19	-0.54	0.22
Home Production TFP	A_p	38.48	1.52	41.38	0.98
Elasticity of Output w.r.t. $\tilde{s}$	γ_1	0.41	0.07	0.59	0.05
Elasticity of Output w.r.t. y	γ_2	0.16	0.12	0.16	0.07
Production Function TFP	A_z	40.50	5.17	42.33	2.28
Female Productivity Wedge	ψ	0.76	0.03	0.84	0.03
Preference shock for Partners (scale)	σ_β^M	0.05	0.01	0.19	0.02

Notes: s.e. denotes standard errors. See Section 7.4 for a description of how these standard errors are computed.

Figure 13: Inequality Changes Over Time: Detailed Decomposition