

MARRIAGE MARKET AND LABOR MARKET SORTING

*PAULA CALVO[†]

ILSE LINDENLAUB[‡]

ANA REYNOSO[§]

July 2023

Abstract

We develop a new 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 equilibrium labor supply as well as marriage and labor market sorting. We then estimate our model using German data to empirically assess the nature of home production, and find that spouses' home hours are *complements*. We investigate to what extent complementarity in home hours drives sorting and inequality. We find that home production complementarity strengthens positive marriage sorting and reduces the gender gap in hours and in labor sorting. This puts significant downward pressure on the gender wage gap and on within-household income inequality, but 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 toward more complementarity—on the evolution of inequality over time.

Keywords. Sorting, Marriage Market, Labor Market, Hours, Household Income Inequality, Gender Wage Gap, Home Production, Technological Change.

*We thank participants at the STLAR Conference at the St. Louis Fed 2019, the Cowles Macroeconomics Conference 2019 at Yale, UPenn, University of Washington St. Louis, Frankfurt University, the H2D2 Research Day at the 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 Andrés, 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, University of Otago, the University of Chile, Cornell, University of Toulouse, University of Western Ontario, Georgetown University, UCLA, Stanford, the Society of Economics of the Household Annual Conference 2021, the Western Economic Association International 96th Annual Conference, the IZA/SOLE Transatlantic Meeting of Labor Economists 2021, the WELAC (LACEA) Workshop on Gender Inequality, the NBER Economics of Culture and Institutions Meeting, the 2022 Cowles Conferences on Labor and Public Economics at Yale, and University of Chicago for useful comments. We thank Hector Chade for very helpful comments on an earlier draft of the paper. The paper also benefitted from discussions with Joseph Altonji, Limor Golan, and Esteban Peralta. We thank Valentina Goetz for outstanding research assistance.

[†]Arizona State University, **Address:** Department of Economics, W.P. Carey School of Business, 501 E Orange St, Tempe, AZ 85283, US. **Email:** paula.calvo@asu.edu

[‡]Yale University, **Address:** Department of Economics, Yale University, 28 Hillhouse Avenue, New Haven, CT 06520, US. **Email:** ilse.lindenlaub@yale.edu

[§]University of Michigan, **Address:** Department of Economics, University of Michigan, 611 Tappan Av., Ann Arbor, MI 48109, US. **Email:** areynoso@umich.edu

1 Introduction

Positive assortative matching—a defining feature of both the labor market and the marriage market—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 shapes inequality has not yet been studied.

This paper shows that sorting in the marriage market and in the labor market are linked by households’ time allocation choices—how time is divided between market work and home production—and examines 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 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 the labor market 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 on spouses’ education in the marriage market and also between workers’ education and jobs’ skill requirements in the labor market. Importantly, there is a gender gap in labor market sorting whereby, 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 the same education allocate 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 (with home hours producing the household’s public good), as well as their private consumption. Last, in the labor market, individuals match with jobs of different productivity, which determines their wages.

The crucial feature of our model is that in the labor market, employers value *both* workers’ skills and hours worked, since hours worked increase individuals’ productivity.¹ Matching between workers and jobs is then based on workers’ *effective* skills—an increasing function of both skills and hours—and jobs’ 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, thereby affecting the value from marriage. Therefore, labor market sorting also impacts marriage market sorting. This interrelation between the two markets and sorting margins is the unique feature of our model—but also renders the problem complex.

We focus on a tractable transferable utility (TU) representation of our model and derive two benchmark equilibria that depend on the model’s primitives. Both equilibria feature positive sorting between workers and jobs in the labor market driven by productive complementarities. However, they differ in household and marriage outcomes depending on the 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 in which spouses are similar in terms of skills and their split between work and home production. The complementarity in home hours is therefore a force toward positive marriage sorting as well as balanced labor supply, labor market sorting, and pay across gender. This leads to a narrow gender wage gap and 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 narrow between-household inequality.

The main 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 must be investigated empirically.

We then study the nature of the home production technology and its role in inequality in the data, both in the cross-section and over time. To do so, we minimally augment our model to capture additional sources of observed heterogeneity while preserving its parsimony and core mechanism. First, we introduce three shocks: marriage taste shocks to allow for mismatch in the marriage market, labor supply shocks to capture time use variation within each couple type, and a random component of workers’ skills to account for mismatch in the labor market. Second, we parameterize our model and allow for gender differences in both home and labor market productivity (the latter can also be interpreted as

¹We base this assumption on our evidence of a positive impact of labor hours on *hourly* wages in the GSOEP, see Figure O.5 and Table O.12, column (3) (both Online Appendix); and also on previous evidence that more labor hours lead to higher productivity—for instance due to reduced coordination costs among co-workers (e.g., [Goldin, 2014](#)).

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 critical 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.

Our main quantitative exercise focuses on Germany over time and investigates 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 using data from the 1990s and compare it with our baseline estimation. Our estimates reveal significant changes in home production over time, with modern Germany being characterized by stronger complementarity in spouses’ home hours and increased relative productivity of men, indicating a switch toward a more ‘progressive’ economy (the monotone equilibrium of our model). These changes in home production technology account for more than 50% of the observed decline in the gender wage gap and for the entire drop in within-household inequality. Half of this drop in gender inequality is due to more gender-balanced productivity at home, but the other half is due to increasingly complementary home production hours. This triggers a decline in the gender gap in both labor hours and labor market sorting and an increase in marriage market sorting—all forces toward more gender equality. 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 and prevented gender gaps from narrowing further.

When isolating the role of sorting, 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 be 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 reduced gender disparities, since it was predominantly driven by women’s improved labor sorting; this helped them catch up with men’s pay.

Given the prominence of home production complementarities in our analysis, we end by providing additional evidence to understand their main sources. Empirically, we find that a substantial boost to spousal complementarities in *childcare* was the main driver behind the increase in aggregate home production complementarities. Consistent with this finding, our model estimation, when implemented separately on samples of couples with and without children, shows that home production complementarities are significantly higher for couples *with children*. Our paper indicates that these changes in how home production is organized within couples had fundamental impacts not only on the marriage market but also on the labor market and, ultimately, on inequality.

Related Literature. This paper contributes to four strands of the literature, as follows.

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, whereby family and fertility choices have a permanent effect on the gender earnings gap even if the wage rate is exogenously fixed (Angelov, Johansson, and Lindahl, 2016; Adda, Dustmann, and Stevens, 2017; Costa Dias, Joyce, and Parodi, 2021; Kleven, Landais, and Sjøgaard, 2019). In our case, the wage rate itself depends on hours worked, which is a key feature of our model. In assuming that hours worked affect workers’ productivity in the market, we follow more closely the literature that documents significant labor market returns to hours (Aronson and French, 2004; Gicheva, 2013; Goldin, 2014; Cortés and Pan, 2019; Bick, Blandin, and Rogerson, 2022).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, 2022) and to sorting into occupations that require different time inputs (Erosa, Fuster, Kambourov, and Rogerson, 2022). Finally, there is literature on the importance of information frictions for gender pay gaps (without considering the marriage market): If employers believe that women have less market attachment than 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 core factor in the gender pay gap. However, in contrast to both the purely empirical and the structural papers we cite, our work takes into account an endogenous marriage market that 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; Greenwood, Guner, Kocharkov, and Santos, 2016; Greenwood, Guner, and Vandenbroucke, 2017; Eika, Mogstad, and Zafar, 2019). We confirm these findings on positive marriage sorting on education in Germany.

Another approach studies marriage market sorting using structural models. Scholars have investigated how premarital investments in education interact with marriage patterns in a static model (Chiappori, Iyigun, and Weiss, 2009) or in a dynamic life-cycle setting (Fernández, Guner, and Knowles, 2005; Chiappori, Costa-Dias, and Meghir, 2018), and how post-marital investments in a partner’s career interact with marriage and divorce (Reynoso, 2022). Further, structural work analyzes 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 unilateral divorce (Fernández and Wong, 2017, Reynoso, 2022);³ or different tax systems (Gayle and Shephard, 2019) affect household

²Bick et al. (2022) find that the 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, which justifies our decision to not allow for non-monotone effects of hours on wages in our model (we do allow for nonlinear effects).

³In an important contribution, Fernández and Wong (2017) study the welfare effects of unilateral divorce in a model of endogenous marriage, divorce, labor supply, and savings; their focus is not on marriage *sorting*. In another 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*.

behavior and the marriage market equilibrium. Finally, in models with *exogenous* marriage sorting, [Fernández and Rogerson \(2001\)](#) analyze the effect of increased marriage sorting on wage inequality, while [Lise and Seitz \(2011\)](#) focus on its effect on between-/within-household consumption inequality; and [Fernández and Wong \(2014\)](#) study the effect of changes in marriage sorting, divorce probabilities, and wage structure on female labor force participation.

As in those papers, marriage market sorting is an important margin in our model. While we treat education as exogenous, we could think of the choice of how many hours to work as an ‘investment’ in individuals’ effective skills. This investment is impacted by marriage sorting while also impacting labor market sorting, which differs from prior work that tends to keep the labor market exogenous.⁴

LABOR MARKET SORTING. A body of literature investigates sorting on the labor market and documents 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, 2020](#); [Lindenlaub and Postel-Vinay, 2023](#)) without taking the marriage market into account. In turn, [Pilossoph and Wee \(2021\)](#) consider spousal joint search on the labor market to explain the marital premium, but take marriage market sorting as given. Our contribution is to examine how the forces that determine who marries whom shape labor market sorting and pay.

INTERPLAY BETWEEN MARRIAGE AND LABOR MARKETS. 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 ([Flabbi, Flinn, and Salazar-Saenz, 2020](#); [Pilossoph and Wee, 2023](#)); changes in wage structure ([Fernández, Guner, and Knowles, 2005](#)); and technological progress in home production ([Greenwood, Guner, Kocharkov, and Santos, 2016](#), [Chiappori, Salanié, and Weiss, 2017](#)) 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 analyzes an equilibrium matching model of both the marriage and the labor market and their interaction. Jointly considering marriage and labor market *sorting* is novel, as is our mechanism regarding how the two sorting margins are linked (i.e., through endogenous labor supply); our finding on the key role of home production complementarities/substitutabilities for shaping equilibrium, both in theory and in the data, is also new.

2 Descriptive Evidence

We first present evidence related to sorting in the marriage market, sorting in the labor market, and the interaction between them. We then highlight the fact that the allocation of hours to paid work and home production is an important link between the two markets.

We use two data sources. The German Socioeconomic Panel ([SOEP, 2019](#)) is a household panel of around 25,000 individuals (including household heads and their spouses), surveyed yearly. It contains

⁴Exceptions are [Fernández and Rogerson \(2001\)](#) (but marriage sorting is kept exogenous) and [Fernández et al. \(2005\)](#) (who endogenize the wages of low- and high-skilled workers, but their model lacks labor market sorting).

detailed information on labor market outcomes and time use. We focus on West Germany, 2010-2016. In turn, the BIBB Employment Survey of 2012 (Hall, Seifer, and Tiemann, 2020 and Rohrbach-Schmidt and Hall, 2020) contains occupational characteristics. Details on the datasets and the construction of our main variables are in Online Appendix OD.1 and OD.2.

MARRIAGE MARKET SORTING. We document positive assortative matching (PAM) on education in the German marriage market, in line with previous evidence (Eika et al., 2019; Greenwood et al., 2016; Greenwood et al., 2017). Table 1 reports marriage market matching frequencies by education for the period 2010-2016 and suggests that almost 60% of individuals marry someone with the same education level (the analysis for 1990-1996 is in Table O.1, Online Appendix OA.1). The correlation between the education levels of spouses—our summary measure of marriage market sorting—equals 0.48.⁵

Table 1: Marriage Matching Frequencies by Education

	Low Education Men	Medium Education Men	High Education Men
Low Education Women	0.16	0.06	0.03
Medium Education Women	0.13	0.25	0.11
High Education Women	0.03	0.05	0.17

Notes: *Low Education* includes either only high school degree or a middle school degree plus basic vocational education (with < 11 years of schooling). *Medium Education* includes any secondary degree plus vocational education (with ≥ 11 years). *High Education* is defined as college or more. We consider an individual’s maximum educational attainment and keep only one observation per couple.

LABOR MARKET SORTING. We measure labor sorting as the correlation between workers’ and jobs’ attributes, where a job is defined by the occupation of the individual (note that the GSOEP does not have firm identifiers). The match-relevant characteristic of workers in the labor market is ‘education’. In turn, the match-relevant attribute of jobs is their ‘task complexity’, constructed from information on the task requirements of each occupation.⁶ The correlation between workers’ years of education and jobs’ task complexity is 0.62, which indicates positive assortative matching on the labor market.


Figure 1 (left) plots the fitted labor market matching function (job attribute as a function of worker characteristic) by gender (solid lines).⁷ Conditional on employment, both men and women are positively sorted in the labor market, indicated by the positive slope of the matching function. However, men are ‘better’ matched: For a given education level, men are on average matched to more demanding jobs. Moreover, the correlation of worker and job attributes differs by gender: It is 0.64 for men and 0.62 for women. A simple empirical check reveals that this gender gap in labor market sorting can account for

⁵As discussed in Chiappori, Dias, and Meghir (2020), *correlation* is a proper measures of sorting. Eika et al. (2019) propose an alternative sorting measure, based on the likelihood that we observe couples with similar levels of education compared with what this frequency would be under random matching. This measure equals 1.77 in 2010-2016: Individuals are 77% more likely to marry someone with the same education, relative to random matching. During 1990-1996, this sorting measure is 1.62, which suggests increased sorting over time (Tables O.2 and O.3 in Online Appendix OA.1).

⁶In Online Appendix OE.4, we compute the task complexity measure for each occupation based on its task content. An occupational type y is the occupation’s rank in the task complexity distribution, where $y \in [0, 1]$. This measure captures the cognitive content of occupations. An advantage over commonly used measures that rank occupations based on wages is that our measure is less impacted by the endogenous selection of workers into those occupations.

⁷All figures in this section show a quadratic fit. For instance, in Figure 1 (left), we plot the prediction for y from a regression of y on x and x^2 with 95% confidence interval in gray, and similarly for the other figures.

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



a substantial part (41%) of the gender wage gap conditional on education.⁸

LABOR MARKET SORTING AND MARRIAGE MARKET SORTING. Next, we assess the relationship between labor market and marriage market sorting (Figure 1, right). For graphical illustration, we measure marriage market sorting by the difference between the years of education of an individual and those of their partner, with ‘zero’ indicating maximum sorting (depicted by the green vertical line). As above, we measure labor market sorting as the correlation between workers’ years of education and the task complexity of their job. The striking—and we believe novel—feature is that labor market sorting is maximized when marriage market sorting is maximized for both men (blue) and women (red).

To control for covariates (especially education) that could affect the link between both sorting margins, and also to alleviate potential attenuation bias due to the changing variability of education across marriage market sorting bins, we complement this graphical analysis with a regression framework:

$$\text{TaskComplexity}_{its} = \beta_0 + \beta_1 \text{Educ}_i + \beta_2 \text{Educ}_i \times \text{PAM}_{it} + \beta_3 \text{PAM}_{it} + X_{it} \mathbf{\Gamma} + \delta_t + \delta_s + \epsilon_{its}, \quad (1)$$

where $\text{TaskComplexity}_{its}$ is the percentile rank of individual i ’s occupation in the task complexity distribution in year t and state s , defined in Online Appendix OE.4. Educ_i is the highest education level attained by an individual, defined as in Table 1. PAM_{it} is an indicator that takes value 1 when both individual i and their partner have the same level of education in year t . Finally, vector X_{it} includes demographic controls, and δ_t and δ_s capture year and state fixed effects.

We implement (1) for men and women and also for the pooled sample (Table O.4, Online Appendix OA.2). Our results show a positive correlation between individuals’ education and their jobs’ task complexity in all samples (given by a positive and significant β_1), which reaffirms positive labor market sorting. This effect becomes *larger* when individuals are well matched in the marriage market, suggested

⁸We compare the coefficient on ‘male’ in a regression of log hourly wages on education and a male dummy with the coefficient on ‘male’ in a regression that additionally controls for task complexity y in a flexible way. Controlling for task complexity induces a drop in the gender wage gap (conditional on education) from 23.9 to 14.2 log points (41%).

by a positive and significant β_2 . For instance, in the pooled sample (column 3), moving up to the next education level increases the average job attribute by 18.8 percentage points. But for those who are perfectly matched in the marriage market this effect is amplified by 7.1 percentage points. This is consistent with the inverse U-shaped relation between labor and marriage market sorting in Figure 1 (right).

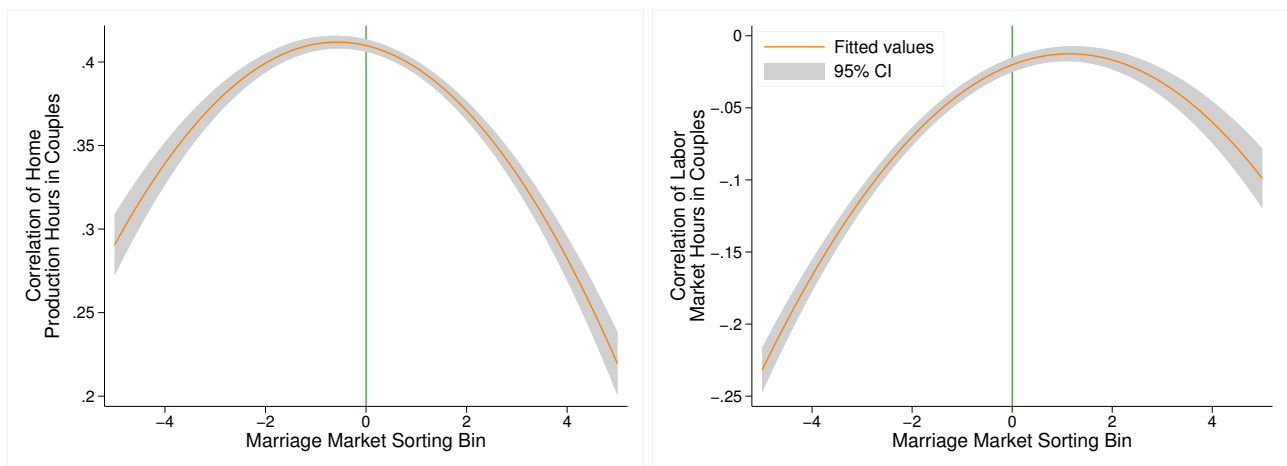
Finally, in Figure O.1 and Table O.5 (Online Appendix OA.2) we replicate the graphical and regression analyses after splitting the sample by education level.⁹ These results suggest that our findings on the relationship between marriage and labor market sorting are not driven by individuals with a specific education level (e.g., by the highly skilled) but hold *within* all broad education levels.

THE ROLE OF HOURS. We now provide evidence on a salient link between the two markets: hours worked on the labor market versus hours spent in home production. First, we show that time allocation is related to partnership status and to marriage market sorting. Second, we document that time allocation choices are also linked to labor market sorting and wages.

As is well documented (Goussé et al., 2017b; Gayle and Shephard, 2019), an individual’s time allocation between ‘work’, ‘home production’, and ‘leisure’ is related to their partnership status. Figure O.4 (Online Appendix OA.4) shows that gender differences in time allocation are small for singles (left panel) but pronounced for couples (right panel). Indeed, women in couples spend 12.5 fewer hours per week in the labor market but about 20 hours more in home production, compared with their male partners. Note that there are no significant differences in leisure by marital status (nor by gender or education).

We also document the relationship between hours and marriage market sorting. Figure 2 plots the correlation between partners’ home production hours (left) and partner’s labor hours (right) against our measure of marriage market sorting (difference in partners’ years of education). Interestingly, both for home production and labor market hours, the correlation is higher when partners are well sorted in the marriage market, as indicated by the inverse U-shape of the hours’ correlation functions.

Figure 2: Time Allocation and Marriage Sorting



A potential concern with Figure 2 is that our marriage market sorting bins pool individuals with dif-

⁹To have variation in education within each subsample in regression (1), we replace *Educ* (level) by *Years of Education*.

ferent education. If hours only depend on own education but do not vary with partner’s education—e.g., if low (high) educated workers always supply low (high) hours independent of the partner’s type—this would generate the inverse U-shape in Figure 2 but be unrelated to spousal complementarities in hours.

We address this point with a regression framework, which allows us to control for the individual’s education and other demographics that can affect hours choices:

$$\text{MaleHours}_{cts} = \alpha_0 + \alpha_1 \text{FemaleHours}_{cts} + X_{ct}\mathbf{\Gamma} + \delta_t + \delta_s + \epsilon_{cts}, \quad (2)$$

where FemaleHours_{cts} and MaleHours_{cts} are either home production or labor market hours of the female and male partners in couple c at time t and state s . The vector X_{ct} includes controls for male education, male age, and the presence of children. Finally, δ_t and δ_s are fixed effects as in (1).

Table O.6 (Online Appendix OA.3.1) shows the results for *home production hours*. A positive and significant coefficient α_1 indicates spousal complementarities in home production (column 1), which are stronger among well-sorted couples (column 2), consistent with Figure 2 (left). In terms of magnitudes, a one hour increase of female home production time is associated with an increase of 0.19 home hours by her male partner. The effect on male hours is larger (0.21) if both partners have the same education.

Table O.7 (Online Appendix OA.3.2) shows the results for *labor market hours*. Our coefficient of interest (α_1) is close to zero for both the whole sample and same-educated spouses (columns 1 and 2). Different factors may bias α_1 downward: First, reported hours are subject to measurement error leading to attenuation bias. Second, both partners’ hours may be driven by omitted factors that correlate positively with male but negatively with female labor hours (e.g., a promotion into a ‘greedy’ job that induces the male partner to work more, especially under reduced flexibility, results in a reduction of hours of the female partner, since she needs to be on call for childcare—a point made by Goldin, 2021).

To address these concerns, we leverage a policy change that induced (exogenous) variation in childcare availability in Germany across time and space, which allows us to instrument for female labor hours in regression (2). Details on the instrument and identification assumptions are in Online Appendix OA.3.2.¹⁰ The IV results are reported in Table O.7.¹¹ Once we instrument for female labor hours, we find a positive and significant impact of wives’ labor hours on husbands’ labor hours (column 3), which indicates complementarities between spouses’ time inputs. The effect is even larger for well-sorted couples in the marriage market (column 4), in line with Figure 2 (right). Indeed, a one hour increase of female work time induces her male partner to work 0.35 hours more (0.51 in same-educated couples).

Finally, for robustness, we repeat the analysis after splitting our sample into three groups based on the education level of the male partner. Figures O.2 and O.3, as well as Tables O.8 and O.9, suggest that our results on partners’ time complementarities tend to hold, also conditional on male education

¹⁰Clearly, these concerns (omitted variables and attenuation bias) are also present in the regression for home production hours. Unfortunately, our IV for labor hours is not valid for home hours since we find that the presence of a small child and childcare availability both have a *direct* impact on male home production, which violates the exclusion restriction. However, based on the expected (downward) bias, our estimate of home production complementarities may be a lower bound.

¹¹In column 5 of Table O.7, we report the first stage corresponding to the IV regression in column 3 (F-statistic=25).

(see Table O.8 for home production and columns 3-4 of Table O.9 for labor hours, where couples with medium-educated men—by far the largest subsample—drive the result).

Besides relating time allocation and *marriage market* outcomes, we stress that the time split between labor market and home production is connected to *labor market* outcomes. First, we document a sizable hourly wage penalty for working part-time (Figure O.5, left panel, in Online Appendix OA.5). In particular, while full-time women have a wage penalty of 14.7 percentage points relative to full-time men, the female part-time wage penalty is 24 percentage points (see [Aaronson and French, 2004](#); [Goldin, 2014](#); and [Bick, Blandin, and Rogerson, 2022](#) for related evidence from the US). Moreover, while less than 10% of employed men work part-time, more than 50% of employed women do so (Figure O.5, right) and are thus particularly affected by these wage penalties.

However, these effects of hours on wages cannot be interpreted as causal due to various endogeneity concerns. Unobserved time-invariant heterogeneity, such as latent productivity or taste for certain work arrangements, may affect both hours and wages. Also, time-varying omitted factors, such as productivity and health shocks, or a change in non-pecuniary job benefits, may affect both work hours and wages. To address these concerns, we identify the effect of hours on the hourly wage in a panel regression with individual fixed effects, in which we instrument for labor hours through the partner’s labor hours (see Online Appendix OE.3.3 for details). We again find a significant wage penalty for not working full-time: An increase from 30 to 40 weekly hours raises the hourly wage by around 4%, which suggests that hours are a productive attribute in the labor market, a point also highlighted by [Goldin \(2014\)](#).¹²

Finally, the underlying force behind the effect of hours on pay may in part be the positive impact of hours on labor market *sorting*. Indeed, accounting for gender differences in hours worked considerably shrinks the discrepancy between the male and female matching functions. We show this in Figure 1, left panel, where solid lines represent the matching functions by gender and dashed lines plot the residualized matching functions after partialling out hours worked.¹³ Still, a small gender gap in labor market sorting persists after controlling for hours worked, which must be accounted for by other factors.

SUMMARY. We highlight three sets of facts. First, we document PAM in both the labor and the marriage market. However, in the labor market, men are ‘better’ matched than women. Second, labor market sorting is maximized when marriage market sorting is. Third, the split between hours worked in the labor market versus home production is a salient link between the two markets: Time allocation choices depend on marriage market sorting but also impact labor market sorting and the hourly wage. Motivated by these facts, we now build a model with endogenous labor *and* marriage markets, linked through households’ labor supply choices. We return to these facts when validating our model below.

¹²This estimated effect in our sample of men and women in Germany is smaller (see column 3 of Table O.12) but comparable to effects estimated on US data: [Aaronson and French \(2004\)](#) also use a panel regression with fixed effects and an IV for hours. They find that increasing hours from 20 to 40 per week increases the hourly wage by 25%. [Bick et al. \(2022\)](#) (who focus on men) find that increasing hours from 30 to 40 per week increases the hourly wage by 11%.

¹³To partial out hours (here and in Section 4.5), we first regress workers’ job attributes on their labor hours and then we regress the residual on years of education. Note that the gender gap in labor market sorting also narrows when we condition on full-time work, but not as much as when controlling for hours as in Figure 1.

3 The Model

We first lay out the environment and decisions, and define equilibrium. Then, we present our results on how this model captures the empirical facts presented in Section 2 in a qualitative way.

3.1 Environment

The model is static (and lasts for one unit of time). There are two types of agents: individuals and firms/jobs (where we use the terms ‘firms’ and ‘jobs’ interchangeably; the empirical counterpart is ‘occupations’). There is a measure one of firms. Firms have productivity attribute $y \in \mathcal{Y} = [\underline{y}, \bar{y}]$, distributed according to a continuously differentiable cumulative distribution function (cdf) G , with density $g > 0$. 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: Women’s skills are denoted by $x_f \in \mathcal{X}_f = [0, \bar{x}_f]$, distributed with a continuously differentiable cdf N_f with density $n_f > 0$. Analogously, men have skills $x_m \in \mathcal{X}_m = [0, \bar{x}_m]$, distributed according to the continuously differentiable cdf N_m with density $n_m > 0$.

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, our novelty is that not only do skills matter for output but also hours worked, which will be chosen optimally by each couple. Thus, heterogeneous firms match with workers’ *effective* skills—a combination of skills and labor hours. Each individual has a fixed time budget, which we normalize to one. It can be allocated to paid work in the labor market, denoted by $h_i, i \in \{f, m\}$, or non-paid work at home toward the production of a public good, $\ell_i = 1 - h_i$, where $h_i = 0$ captures non-participation in the market.¹⁴ By increasing labor hours, each individual ‘invests’ in his/her effective skill $\tilde{x} := e(x, h)$, $\tilde{x} \in \tilde{\mathcal{X}}$, with endogenous cdf $\tilde{N}(t) := \mathbb{P}[\tilde{x} \leq t] = \frac{1}{2}\mathbb{P}[\tilde{x}_f \leq t] + \frac{1}{2}\mathbb{P}[\tilde{x}_m \leq t]$. We assume that e is twice continuously differentiable, strictly increasing in each argument, strictly supermodular in (x, h) , and $e(0, h) = 0$ for all h . Thus, putting in more labor hours is *as if* the worker is more skilled. Our assumption—that not 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 of a homogeneous final good 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 (\tilde{x}, y) . Because z depends on the effective worker type, \tilde{x} , which in turn depends on labor hours, hours are a productive *input* to labor market production.¹⁵ Match output $z(\tilde{x}, y)$ is split into wages and profits (see below), where worker \tilde{x} uses that wage to finance private consumption c of the final good and, similarly, firm y uses profits to pay for its consumption.

The public good production function is denoted by p . It takes as inputs each couple’s hours at home, so that $p(\ell_m, \ell_f)$ is the public good produced by a couple working $(\ell_m, \ell_f) = (1 - h_m, 1 - h_f)$ in home

¹⁴We abstract from leisure, since we observe only small differences in leisure across gender, marital status, and education.

¹⁵We base this assumption on our evidence (Online Appendix OE.3.3) that more labor hours lead to higher productivity and hourly pay (see also Aaronson and French, 2004; Goldin, 2014; Cortés and Pan, 2019; Bick et al., 2022).

production.¹⁶ We assume that p is twice continuously differentiable, strictly increasing and concave, with standard Inada conditions (i.e., $\lim_{h_i \rightarrow 0} p_{\ell_i}(1 - h_m, 1 - h_f) = 0$ and $\lim_{h_i \rightarrow 1} p_{\ell_i}(1 - h_m, 1 - h_f) = \infty$).¹⁷

The utility function of an individual is denoted 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 C^2 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 the marriage market—are competitive (full information, no search frictions, and price-taking behavior) and there is no risk. The two markets and sorting choices therein are linked through labor supply choices, which can be interpreted as a pre-labor market and post-marriage market continuous investment in ‘effective’ skills. This link is the crucial element of our model.

3.2 Decisions

Agents make three decisions, summarized in Figure 3. In the *marriage market stage*, men and women choose their partner to maximize their value of marriage. The outcome is a marriage market matching function that matches each woman x_f to some man x_m (or singlehood) and a market-clearing price. In the *household decision problem*, each matched couple chooses private consumption and allocates their hours to labor market work and home production, under anticipation of the labor market outcomes (matching and wages). This stage yields both private consumption and time allocation choices (and hence public consumption) and pins down individuals’ effective types \tilde{x} . In 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. We now go into detail.

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

Stage:	<i>Marriage Market</i>	<i>Household</i>	<i>Labor Market</i>
Allocations:	x_m matches with $x_f \rightarrow \eta(x_f)$	$c_f, c_m, h_m, h_f \rightarrow p, \tilde{x}_f, \tilde{x}_m$	\tilde{x}_i matches with $y \rightarrow \mu(x_i)$
Prices:	$v(x_f)$		$w(\tilde{x}_f), w(\tilde{x}_m)$

LABOR MARKET. Given the distribution of effective types \tilde{N} , derived from the marriage and household decisions, and given the wage function $w : \tilde{\mathcal{X}} \rightarrow \mathbb{R}_+$, a firm with productivity y chooses the *effective* worker type that maximizes its profits:

$$\max_{\tilde{x}} z(\tilde{x}, y) - w(\tilde{x}). \quad (3)$$

¹⁶Our model can handle more general home production functions in which part of the public good is purchased using wages. But given that (i) in detailed time-use data (German Time Use Survey), we find limited outsourcing of home production tasks (~ 2 h per week) and (ii) we do not observe purchased public goods in our main dataset used for estimation (GSOEP), we interpret spouses’ home production time as being *net* of what was purchased, in both the data and model.

¹⁷We will denote the *partial* derivatives of some generic function $f(x, y)$ using subscripts; e.g., 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 $(f)_x$.

This problem, along with market clearing, pins down the labor market matching function $\mu : \tilde{\mathcal{X}} \rightarrow \mathcal{Y}$, which maps workers' effective skills to firm types.¹⁸ Matching function μ depends on the complementarities of (\tilde{x}, y) in z . Importantly, it also 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, $\tilde{\mathcal{X}} = [0, \bar{x}]$ —as will be the case in Section 3.4 below—then the first-order condition of problem (3), which gives a differential equation for w , pins down the wage function as

$$w(\tilde{x}) = w_0 + \int_0^{\tilde{x}} z_{\tilde{x}}(t, \mu(t)) dt, \quad (4)$$

where w is the wage per unit of time and w_0 is a constant of integration.

HOUSEHOLD PROBLEM. Consider a couple (x_f, x_m) who takes w from the labor market as given. This couple solves the following *cooperative household problem*. One partner (here wlog the man) 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 hours allocation:

$$\begin{aligned} \max_{c_m, c_f, h_m, h_f} \quad & u(c_m, p(1 - h_m, 1 - h_f)) \\ \text{s.t.} \quad & c_m + c_f - w(\tilde{x}_m) - w(\tilde{x}_f) \leq 0 \\ & u(c_f, p(1 - h_m, 1 - h_f)) \geq \bar{v}, \\ & 0 \leq h_i \leq 1, i = \{f, m\} \end{aligned} \quad (5)$$

where at this stage \bar{v} is taken as a parameter by each household (but it will be a function of female skills and determined in the marriage market stage below). When solved for all feasible $\bar{v} \in [0, \bar{v}_{max}(x_f, x_m)]$ (where $\bar{v}_{max}(x_f, x_m)$ is the maximum that x_f can obtain when matched with x_m), problem (5) traces out the household's Pareto utility frontier. The solution to this problem yields, for each partner in a couple (x_m, x_f) and for a given utility \bar{v} , both private consumption $c_i(x_m, x_f, \bar{v})$ and labor hours $h_i(x_m, x_f, \bar{v})$ with $i = \{f, m\}$ (and therefore the couple's output of the public good, $p(1 - h_m, 1 - h_f)$).

MARRIAGE MARKET. Anticipating the solution to the household problem (h_f, h_m, c_f, c_m) for each *potential* couple, as well as taking the wage function w and transfer function v as given, the value of man x_m from marrying woman x_f is given by the value of household problem (5),

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

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

$$\max_{x_f} \Phi(x_m, x_f, v(x_f)). \quad (6)$$

¹⁸Since we focus on monotone matching below, we restrict attention to pure matching, given by a function μ .

The FOC of this problem (which gives a differential equation for v), together with marriage market clearing, determines the marriage matching function $\eta : \mathcal{X}_f \rightarrow \mathcal{X}_m$, mapping female skills to male skills in a measure-preserving way. It also determines a transfer (in utils) function $v : \mathcal{X}_f \rightarrow \mathbb{R}_+$ that supports allocation η , 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 (x_m, x_f) 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 here if the value of marriage Φ is positive for all potential couples.

3.3 Equilibrium

We now formally define the equilibrium of our model.¹⁹

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

1. *given $(\eta, v, h_m, h_f, \tilde{N})$, the pair (w, μ) 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 because it captures several observed features of the interaction between marriage and labor markets from Section 2.

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 .*

A monotone equilibrium has three additional properties associated with the three stages of this model. Most importantly, there is positive sorting in each market and labor hours are increasing in own and partner's 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 , which implies that γ_f can be inverted, and similarly for men with $\gamma_m(x_m)$. As a result, the endogenous cdf of effective types can be written as:²⁰

$$\tilde{N}(t) = \frac{1}{2}N_f(\gamma_f^{-1}(t)) + \frac{1}{2}N_m(\gamma_m^{-1}(t)),$$

which depends on marriage outcomes through (γ_f, γ_m) . This highlights an important point. The equilibrium hours function, h_f , not only depends on a woman's own skill, x_f , but also on marriage market outcomes: the skill of her partner, $\eta(x_f)$, as well as the transfer $v(x_f)$; and similarly regarding the factors that impact the male hours function h_m . Thus, labor supply choices form the link between the marriage

¹⁹In equilibrium, labor market output z equals profit plus wage *in each match*, so the goods market clears at the match level and, by implication, also economy-wide, which is why we do not specify goods' market clearing explicitly.

²⁰To see this, observe that the probability that $\tilde{x} \leq t$ is

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

market (they are determined by the household and depend on who marries whom, captured by 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 theoretically challenging, as we must simultaneously equilibrate two intertwined matching markets, which are related through time-allocation choices.

3.4 Conditions for Monotone Equilibrium

We show how the model primitives shape equilibrium and, specifically, how complementarities between spouses' hours in the home production function give rise to a monotone equilibrium.

To gain tractability in this equilibrium problem of interconnected markets, we impose two restrictions. (i) We focus on equilibria that satisfy some basic regularity properties (*Regular Equilibrium*), summarized in Definition O1 of Online Appendix OB.2, which can be justified in terms of primitives. In Appendix A, we detail why these regularity assumptions are important for our main result. (ii) We further focus on the quasi-linear class of utility functions, $u(c_i, p) = F(c_i + p)$ (with F strictly increasing),²¹ which yields the *transferable utility* (TU) property.²² Under TU, the household's aggregate demand for private consumption c and public consumption p can be determined irrespective 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, the couple's marital surplus is then also independent of sharing rule v and the matching problem can be solved by maximizing the total value of marriage, independent of how it is shared (as in Shapley and Shubik, 1971 and Becker, 1973).

Proposition 1 (Monotone Equilibrium). *Assume p is strictly supermodular in (ℓ_m, ℓ_f) , and z is strictly supermodular in (\tilde{x}, y) and convex in \tilde{x} for each y . Then any regular equilibrium is monotone.*

The proof, which relies on monotone comparative statics arguments, is in Appendix A but here is the intuition. The crucial condition for the monotone equilibrium is the complementarity of spousal time in home production (supermodular p , $p_{\ell_m \ell_f} > 0$). This gives rise to a 'progressive' way of organizing the household with gender balance in hours as opposed to specialization. In this case, increasing, for instance, female skills not only increases her own labor hours at the cost of fewer home hours but—due to the home production complementarity—also induces her partner to work more in the market and less at home. Thus, partners' hours comove and are thus complementary. This positive correlation of partners' hours within the household is clearly a force toward positive sorting in the marriage market: It makes wages complementary in types, $(w)_{x_m x_f} > 0$ (and thus Φ is supermodular, $\Phi_{x_m x_f} > 0$).²³ This

²¹We here focus on the simplest functional form that yields the TU property and gives us tractable conditions for monotone equilibrium. Online Appendix OB.1 demonstrates that a broad class of utility functions renders the model TU representable and highlights conditions under which the monotone equilibrium still obtains. Thus, focusing on the TU case is not overly restrictive. Moreover, even under imperfectly transferable utility (ITU), home production complementarities remain a crucial force for monotone equilibrium, but the derivations become much more complex (available on request).

²²More generally, the Gorman form yields TU, where i 's utility is given by $u^i(p, c_1, \dots, c_n) = z^i(c_2, \dots, c_n) + k(p)c_1$, which is *linear in at least one private consumption good*, with common coefficient $k(p)$, so that utility can be transferred between partners at a constant rate. See Mazzocco (2007) and Browning et al. (2014).

²³In the quasi-linear class (TU case), the sufficient condition for PAM in the marriage market is simply $\Phi_{x_m x_f} > 0$.

is because the marginal wage return to female skills gets an extra push from a more skilled husband who boosts her labor hours. Finally, positive sorting in the labor market stems from the complementarity between individuals’ effective skills and jobs’ skill requirements (supermodular z).

Indeed, the monotone equilibrium captures several salient features of the data. As we detail in Online Appendix OB.3, the properties of monotone equilibrium are *qualitatively* consistent with our stylized facts from Section 2 (and we accurately replicate our facts in the *quantitative* analysis below). In particular, the monotone equilibrium is in line with positive labor market sorting (Figure 1, left); positive marriage market sorting (Table 1); the fact that labor market sorting is reinforced by marriage market sorting (Figure 1, right); the fact that the hours complementarity between spouses is reinforced by positive marriage sorting (Figure 2); and the fact that controlling for hours worked can close any potential gender gap in labor market sorting (dashed lines in Figure 1, left).

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. It is plausible that in certain settings a different equilibrium arises, in which partners’ hours in home production are substitutable and positive sorting in the marriage market was less pronounced (or sorting was even negative), giving rise to the specialization of household members. We capture this different regime by what we call—with some abuse—a *non-monotone equilibrium*. We define this equilibrium as the monotone one with two differences: (i) *negative* assortative matching (NAM) in the marriage market and (ii) labor hours are *decreasing* in partner’s type.

We show in Proposition O1 (Online Appendix OB.4) that the crucial assumption underlying the non-monotone equilibrium is a *substitutability* in home hours, $p_{\ell_m \ell_f} < 0$. This gives rise to an equilibrium that relies on ‘specialization’: Increasing, for instance, male skills raises his labor hours, while female labor hours go down in response. At the same time, the male partner spends less time in home production, while female home hours increase. This specialization within the household is clearly a force toward NAM in the marriage market, which indeed materializes. The reason is that increasing the partner’s type pushes own labor hours *down*, which hurts own labor market prospects, especially for skilled individuals (here: $(w)_{x_m x_f} < 0$ and thus $\Phi_{x_m x_f} < 0$). Skilled agents then prefer to match with *less* skilled partners.

Thus, complementarity vs. substitutability of home hours shapes equilibrium. Specifically, $p_{\ell_m \ell_f} \leq 0$ determines whether marriage partners match positively and whether their hours, both at home and at work, comove. The monotone equilibrium captures ‘progressive’ societies, while the non-monotone one reflects a ‘traditional’ division of labor.²⁴ To our knowledge, this mechanism, in which home production complementarities are the main driver of marriage *and* labor market outcomes, is new in the literature.

3.5 Bringing the Model to the Data

Our goal is to investigate the nature of home production in the data and to assess the empirical relevance of our model mechanism. To match relevant features of the data in our estimation—particularly,

²⁴See Becker (1985) for a seminal contribution on the specialization of household members.

imperfect sorting and non-participation in both markets, as well as heterogeneous hours choices among otherwise equal couples—we propose three minimal departures from our baseline setting.

First, in order to capture mismatch in the labor market along (x, y) , we augment individuals' education, x , by a 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 above. We assume ν (and thus s) is observed by the agents, but not by the econometrician. In the labor market, the match-relevant attribute of workers is their effective human capital $\tilde{s} := e(s, h)$ (instead of \tilde{x}), whose distribution we denote by \tilde{N}_s . Thus, a firm with productivity y solves $\max_{\tilde{s}} z(\tilde{s}, y) - w(\tilde{s})$ (instead of (3)).

Second, we account for heterogeneity in labor supply (including non-participation) *within* (s_m, s_f) -type couples and *within* s_i -type singles by introducing idiosyncratic labor supply shocks. After the marriage stage, each decision-maker draws a vector of labor supply shocks, one shock for each hours alternative. We denote by δ^{h_i} an agent's idiosyncratic preference for hours alternative h_i , $i \in \{f, m\}$, where hours are discrete elements of choice set \mathcal{H} , $h_i \in \mathcal{H}$. In the household decision stage, the husband (wlog) chooses household labor supply and consumption to maximize his economic utility, as in (5), *plus* labor supply shock, i.e., $u(c_m, p(1 - h_m, 1 - h_f)) + \delta^{h_m}$, subject to the usual constraints (see (5)).

Third, to accommodate imperfect sorting in the marriage market (which is now based on *human capital* s instead of on education x) and to account for singlehood, we introduce an idiosyncratic taste shock for partners' s -types. We denote by β_m^s and β_f^s the 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 to remain single. Thus, individuals in the marriage market value potential partners not only for their impact on the economic joint surplus (as in (6)), but also for their impact on the non-economic surplus (which stems from preference shocks β_f^s or β_m^s). The marriage problem of a man with human capital s_m is therefore: $\max_s \Phi(s_m, s, v(s_m, s)) + \beta_m^s$, where the dependence of v on both partners' types (s_m, s_f) reflects that matching is no longer pure.

We give more details in Appendix B. Importantly, we show in Proposition O2 (Online Appendix OB.5) that under similar conditions as in Proposition 1, the properties of monotone equilibrium hold *on average* in our augmented model. In turn, Online Appendix OC describes the numerical solution of this model. It consists of solving for a fixed-point in the wage function. 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, must coincide with the initially postulated one. 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 to determine whether a particular hours choice is optimal, agents must compare the payoff of this 'investment' with all alternative hours choices. But the competitive wage only determines the price for *equilibrium* hours.²⁵ To obtain the off-equilibrium wages without significantly perturbing the equilibrium wages, we use a *tremble strategy*. We believe the application of trembling to matching markets with investment is new.²⁶

²⁵ A similar issue arises in the pre-match investment problem of Cole, Mailath, and Postlewaite (2001).

²⁶ We postulate that a small fraction of agents are tremblers who make a mistake by choosing off-equilibrium hours. Trembling helps us price all hours alternatives in case the preference shocks for hours do not yield a situation in which all

4 Estimation

We first parameterize our model and show that it is identified. We then estimate it and assess whether partners' home production time is complementary or substitutable in the German data.

4.1 Parameterization

We assume that the labor market production function is given by $z(\tilde{s}, y) = A_z \tilde{s}^{\gamma_1} y^{\gamma_2} + K$, where A_z is a TFP term and (γ_1, γ_2) are the curvature parameters that reflects the elasticity of output with respect to effective human capital and job productivity. In turn, K is a constant that allows for positive output of the least productive match and thus for a positive minimum wage, in line with the data.²⁷

We assume a CES production function of the public good for married couples (M for Married):

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

where A_p is the TFP in home production, θ is the relative female productivity, and ρ 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 (U for Unmarried) is given by $p^U(1 - h_i) = A_p \Theta_i(1 - h_i)$, where $\Theta_i \in \{\theta, 1 - \theta\}$ depending on gender $i = \{f, m\}$.

The utility function of individual i is given by $u(c_i, p) = c_i + p$, where $p \in \{p^M, p^U\}$ for couples and singles. We adjust the private consumption of singles using the McClements equivalence scale to capture the loss of economies of scale in these households (Anyaeabu, 2010).

Human capital follows the functional form $s \propto x + \nu$, i.e., s is proportional to the sum of observed skill x (based on education) and (to us) unobserved productivity ν . The effective human capital functions are

$$\tilde{s}_f = \psi s_f h_f \quad \text{and} \quad \tilde{s}_m = s_m h_m, \quad (7)$$

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 allow for a labor market penalty for women that could reflect, e.g., discrimination or productivity differences.

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

$$\begin{aligned} \beta^s &\sim \text{Type I}(0, \sigma_\beta) && \text{for } s \in \{\mathcal{S} \cup \emptyset\} \\ \delta^{h^t} &\sim \text{Type I}(0, \sigma_\delta) && \text{for } h^t \in \mathcal{H} \text{ and } t \in \{M, U\}. \end{aligned}$$

Moreover, we specify the labor supply shocks as

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

alternatives are chosen in equilibrium. This is important during estimation, in which different levels of the scale parameter of the hours shock distribution are evaluated, including those that would not induce agents to choose all hours alternatives.

²⁷We assume that K is not shared between workers and firms but accrues to the worker. If $K > 0$ ($K = 0$), then the least productive labor market match $(\tilde{s}, \underline{y})$ renders a positive (zero) wage; see wage equation (4) with $w_0 > 0$ ($w_0 = 0$).

That is, when making hours choices, a decision-making unit—either a married or a single household—draws *only one* labor supply shock for their time allocation, δ^{ht} , which is extreme-value distributed.^{28,29}

4.2 Identification

We need to identify 10 parameters and two distributions. We group these objects into 5 categories and discuss their identification group-wise. We have parameters pertaining to the home production function (θ, ρ, A_p) , the labor market production function $(\gamma_1, \gamma_2, A_z, K)$, labor supply and marriage preference shock distributions $(\sigma_\delta, \sigma_\beta)$, and a labor productivity wedge (ψ) . Finally, we have the distributions of worker human capital and job productivity (N_s, G) . 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. We provide formal identification arguments in Appendix C and summarize the logic here.

The distributions of worker and job heterogeneity, (N_s, G) , will be identified directly from the empirical distributions, which allows us to assign a human capital level s (productivity y) to each worker (job).

The home production function p (embodying (θ, ρ, A_p)) is identified from choice probabilities for home hours, which have a tractable form under the assumption that labor supply shocks are type-I extreme-value distributed. The resulting relative choice probabilities for all hours combinations of a couple with given human capital types (s_m, s_f) are impacted by the couples' wages (observed), the scale of the labor supply shock (σ_δ) , identified below), and the home production output (p) , identifying p .

The labor market production function, and thus $(\gamma_1, \gamma_2, A_z, K)$, is identified from wage data. We follow arguments from the literature on the identification of hedonic models (Ekeland, Heckman, and Nesheim, 2004) and take advantage of the tight link between wages and the marginal product (and thus technology) in our competitive environment. In turn, the constant in the production function (K) can be identified from the minimum observed hourly wage.

The pair $(\sigma_\delta, \sigma_\beta)$ associated with our shock distributions is identified as follows. In the absence of labor supply shocks $(\sigma_\delta = 0)$, any two couples of the same type (s_m, s_f) 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 $(\sigma_\beta = 0)$, the model would produce perfect assortative matching in the marriage market with

²⁸Thus for couples, we assume that the *sum* of their shocks is extreme-value distributed. We make this adjustment to the standard setting, in which *each* individual 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. Gayle and Shephard (2019) follow a similar approach in their numerical solution (see their footnote 24), in which they assume households draw one shock for each of the couple's joint hours combinations.

²⁹Note that as is, our parsimonious model (featuring no couple-/single-specific parameters) would have difficulty generating enough singles because the economic value of staying single is relatively small. Allowing for different scales of the shock distributions for the alternatives 'single' and 'married', however, would give rise to a nested logit problem with degenerate (single) nest, which is associated with known identification issues for the scale of the degenerate nest (Hunt, 2000). This is why we calibrate an adjustment to the single value outside of the main estimation. See Appendix D.1 for details.

$\text{corr}(s_m, s_f) = 1$. The extent of marriage market mismatch identifies the scale parameter of preference shocks for partners, σ_β . Note that the standard result in the literature, whereby this 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 results for $(\sigma_\delta, \sigma_\beta)$ rely on the extreme-value assumption of the shock distributions, yielding tractable choice probabilities.

The productivity or discrimination wedge of women, ψ , is identified from the hourly gender wage gap *conditional on hours and s-type*. If there is 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 $\psi \neq 1$. In sum:

Proposition 2 (Identification). *Under the assumed functional forms, the model parameters are identified.*

4.3 The Data

We again use data from the GSOEP combined with information on occupational characteristics from the BIBB. One challenge with the estimation is to bridge our static model with life-cycle features of the data. To deal with this, we construct a dataset that features ‘typical’ outcomes of individuals during the period we observe them. Specifically, we define for each individual the typical occupation (based on a combination of tenure and job-ladder features), their typical labor hours and typical wage for that occupation, and their typical home hours while holding that occupation, as well as the typical marital status. Our baseline analysis focuses on a restricted time period (2010-2016) and age group (22-55). This allows us to assess the typical outcomes in only one life-cycle stage during the prime working age.

In line with our model, we only consider those individuals who are married/cohabiting or single. Our final sample contains 3,857 individuals living in West Germany, around 80% of whom are in couples. In Online Appendix OD.3, we provide details on the sample and variable construction.

4.4 Estimation: Strategy and Results

We propose a two-step estimation procedure. The first step estimates worker and job heterogeneity *outside* of the model. In a second step, given the worker and job distributions, we estimate the remaining parameters *within* the model.

4.4.1 First Step: Calibration Outside the Model

The main objects we estimate in this first step are worker human capital types s (i.e., a combination of (x, ν)) and job productivity types y . Except for x (education), these types are not directly observed.

ESTIMATION OF WORKER TYPES. We make use of the longitudinal structure of the GSOEP to estimate workers’ unobserved heterogeneity ν from individual fixed effects in a panel wage regression.

Estimating unobserved heterogeneity outside of the model by making use of the dynamic nature of the data allows us to circumvent some known identification hurdles in static matching models with unobserved types.³⁰ For implementation, we interpret the unit of time in our model as 1 hour, so that it features wages per hour (or hourly earnings). We then specify and estimate an empirical model for *hourly* wages as a function of effective types (which in turn are a function of education x , ability ν , and hours worked h). Based on individuals’ estimated fixed effects and their education level (low, medium, and high, as above), we classify workers into six distinct human capital types s_i , whose empirical cdf is our estimate for their human capital distribution N_s .

We address two challenges in implementing the panel wage regression. First, we use an instrumental variable (IV) approach to account for the endogeneity of labor hours. Second, we apply a Heckman selection correction (Heckman, 1979) to account for non-random selection in labor force participation. We provide details in Online Appendix OE.3. The estimation results of the panel wage regression are in Table O.12 and the estimated skill distribution in Table O.13.

ESTIMATION OF JOB TYPES. The empirical counterpart of our model’s firms/jobs are occupations. As in Section 2, we measure occupations’ productivity y from data on their task complexity. Our main dataset is the BIBB (comparable to the O*NET in the US), which contains extensive information on the tasks performed in each occupation. We apply principal component analysis (PCA) to collapse the multiple task dimensions to a single one, where we use the (normalized) first principal component as our one-dimensional occupational characteristic y . The loadings in the PCA indicate that our measure of y can be interpreted as the ‘cognitive task content/productivity’ of an occupation. Online Appendix OE.4 provides the details of this approach.

EXTERNALLY CALIBRATED PARAMETERS. Outside of the main estimation, we calibrate both the constant of the labor market production function K as well as an adjustment to the value of singles that allows us to match the extensive margin of marriage while circumventing known identification issues (also see footnote 29). Appendix D.1 contains the details and Table A.1 the calibrated parameters.

4.4.2 Second Step: Internal Estimation

There are nine remaining parameters of the model, $\Lambda \equiv (\theta, \rho, A_p, \gamma_1, \gamma_2, A_z, \psi, \sigma_\delta, \sigma_\beta)$. They are disciplined by 16 moments whose choice is guided by our identification arguments (Proposition 2). To estimate the home production function, we use five moments related to the division of labor and to the complementarity of hours within households (the 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). We specifically emphasize the estimation

³⁰We are not aware of an identification result for unobserved heterogeneity in this type of static (hedonic) matching model of the labor market. The literature on the identification of hedonic matching models with unobserved heterogeneity commonly maintains the following two assumptions: (i) the distribution of unobserved heterogeneity is known and (ii) the part of the production function that involves unobserved heterogeneity is also known. Under these assumptions, the part of technology that involves observed heterogeneity can be identified (e.g., Chernozhukov, Galichon, Henry, and Pass, 2021). We cannot impose these assumptions here, since they postulate that the very objects we aim to identify are known.

of ρ —a central parameter in our analysis that captures the extent of home production complementarities.³¹ The main moment informing this complementarity parameter is the correlation of spouses’ home production hours. We show in Lemma 1 (Appendix D.2.1) that a (positive) affiliation between spouses’ hours in the data—which implies a positive correlation—can only be rationalized by home production complementarities in the model. This justifies why we chose the hours correlation to pin down ρ .

In turn, to estimate the labor market production function, we use four moments related to the hourly wage distribution (its mean, variance, and the 90-10 and 90-50 percentiles). To estimate the scale of the marriage shock, we use the extent of marriage market sorting (the correlation of spouses’ human capital types). To pin down 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 two human capital types). Finally, we pin down the female labor wedge ψ with two moments related to the gender wage gap conditional on (s, h) . The construction of these moments is described in Online Appendix OE.1. To estimate these parameters, we apply the simulated method of moments (McFadden, 1989; Pakes and Pollard, 1989); see Online Appendix OE.2.

The estimated parameters (including their standard errors) are in Table 2.³² Regarding the home production function, our estimates indicate that spouses’ hours at home (and therefore in the labor market) are *complements* with $\rho = -0.54$. This implies an elasticity of substitution of spousal home production inputs of 0.65, which pushes the model toward a *monotone equilibrium*. The main data moment that calls for a negative ρ is the strong positive correlation of spouses’ home hours. We further find that women are significantly more productive at home than men ($\theta = 0.82$). The large differences in labor force participation and full-time work across genders call for this high female productivity at home.

In terms of labor market production, our estimates indicate that it is concave in both worker and job productivity ($\gamma_1 < 1, \gamma_2 < 1$).³³ Labor market TFP, A_z , is estimated to be higher than home production TFP, A_p . Finally, the empirical gender wage gap conditional on hours and human capital calls for a female productivity/discrimination wedge, which we estimate as $\psi = 0.85$. This implies that for any given type and choice of hours, women’s effective human capital is 15% lower than that of men.

The last column of Table 2 presents our sensitivity analysis (Andrews, Gentzkow, and Shapiro, 2017). We report the three most important moments for each parameter in estimation.³⁴ The sensitivity

³¹In our equilibrium context, more standard partial equilibrium estimation approaches for ρ based on the ratio of FOCs of the household problem using appropriate IVs for wages (e.g., Rupert, Rogerson, and Wright, 1995; Aguiar and Hurst, 2007; or Moschini, Forthcoming) are difficult to implement. The reasons are that (i) these FOCs do not hold in our quantitative setting since hours are discrete, subject to preference shocks, and do not always satisfy an interior solution; and (ii) even if we used the ratio of FOCs from the *baseline* model, its log-linearized version is subject to a difficult endogeneity problem since not only the dependent variable but also several independent variables are a function of the couple’s labor hours. This is because marginal wage returns are a function of hours in our equilibrium labor market.

³²The covariance matrix of the estimator is computed as matrix $Var = [D'_m \mathcal{V} D_m]^{-1} D'_m \mathcal{V} C \mathcal{V}' D_m [D'_m \mathcal{V} D_m]^{-1}$, where D_m is the 16×10 matrix of the partial derivative of moment conditions with respect to each parameter evaluated at the estimates; C is the covariance matrix of the data moments; and \mathcal{V} is the weighting matrix used in estimation.

³³The convexity of z is only sufficient but not necessary for monotone equilibrium.

³⁴We compute the sensitivity of each parameter to the moments as $Sensitivity = [D'_m \mathcal{V} D_m]^{-1} D'_m \mathcal{V}$, defined by Andrews et al. (2017); see footnote 32 for notation. Since our moments have different scales, we multiply the columns of our *Sensitivity* matrix by the standard deviation of the corresponding empirical moment.

analysis is in line with our identification arguments. For example, the correlation of spouses' hours (moment $M5$) is an important moment that disciplines the home production complementarities, ρ ; or the female productivity wedge, ψ , is strongly related to the within-type gender wage gaps, $M11$.

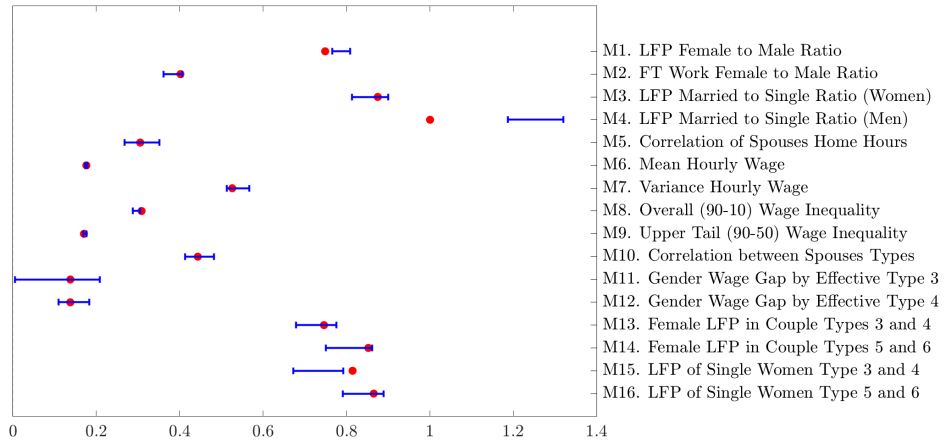
Figure 4 summarizes the fit between model and data moments (see also Table A.2, Appendix D.2). We plot our 16 model moments and their corresponding data confidence intervals. Our model achieves a good fit with the data despite its parsimony, with nearly all model moments lying in the confidence interval of their data moments.³⁵

Table 2: Estimated Parameters

Parameter	Estimate	s.e.	Top-3 Sensitivity Moments
Female Relative Productivity in Home Production, θ	0.82	0.06	M1, M5, M11
Complementarity Parameter in Home Production, ρ	-0.54	0.20	M3, M5, M11
Home Production TFP, A_p	38.33	3.46	M1, M11, M13
Elasticity of Output w.r.t. \tilde{s} , γ_1	0.63	0.05	M7, M12, M13
Elasticity of Output w.r.t. y , γ_2	0.18	0.05	M1, M13, M14
Production Function TFP, A_z	42.00	2.29	M1, M5, M11
Female Productivity Wedge, ψ	0.85	0.02	M7, M9, M11
Labor Supply Shock (scale), σ_δ	7.57	0.93	M1, M11, M15
Preference Shock for Partners (scale), σ_β	0.11	0.01	M10, M11, M14

Notes: s.e. denotes standard errors. Column *Top-3 Sensitivity Moments* reports the three most important moments for each parameter in estimation based on our sensitivity measure (see footnote 34). On average, these three moments jointly explain 48% of the total sensitivity. $M1, \dots, M16$ denote our 16 targeted moments (defined in Figure 4 and Table A.2, Appendix D.2).

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



Notes: The red dots indicate the level of the model moments while the blue bars are their corresponding data 90%-confidence intervals, computed from a bootstrap sample. We rescaled moments $M6 - M9$ to be able to plot all moments in the same graph.

³⁵The only moment our model has difficulty matching is the ratio of LPF of married to single men ($M4$), which, as is well-known, in the data is larger than one. This could, for instance, be due to selection (more productive men marry); the fact that marriage allows men to accumulate more human capital; the fact that married men are healthier; or cultural pressure that pushes married men into the breadwinner role. None of these forces is present in our model.

4.5 Model Validation

Apart from fitting the aggregate moments targeted in estimation, our model reproduces rich, *untargeted* features of the data, documented in Section 2: the relation between marriage and labor market sorting and how hours form the link between them.

MARRIAGE MARKET SORTING. Table 3 displays the marriage market matching frequencies by our three education levels in the model and data. In the estimation, we only targeted the overall correlation of couples’ *human capital* types (i.e., *s*-types), since *s* is the relevant matching characteristic in the marriage market of our model. But we did not target marital matching on *education*, *x*, and especially not the detailed matching frequencies. Nevertheless, the model matches well these observed marriage frequencies: 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; this is another feature of the data we did not target.

Table 3: Untargeted Moments: Marriage Matching Frequencies—Model and (Data)

	Low Educ Men	Medium Educ Men	High Educ Men	Single Women
Low Educ Women	0.0693 (0.1306)	0.0792 (0.0667)	0.0363 (0.0188)	0.0297 (0.0391)
Medium Educ Women	0.1056 (0.0895)	0.1122 (0.1850)	0.0891 (0.0718)	0.0990 (0.0709)
High Educ Women	0.0363 (0.0116)	0.0495 (0.0328)	0.0924 (0.0846)	0.0363 (0.0453)
Single Men	0.0660 (0.0563)	0.0693 (0.0595)	0.0297 (0.0374)	

Notes: *Low Educ* includes either only high school degree or a middle school degree plus basic vocational education with < 11 years of schooling. *Medium Educ* includes any secondary degree plus vocational education with ≥ 11 years of schooling. *High Educ* is defined as college or more. We display model frequencies, with data frequencies in parentheses.

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 features positive labor market sorting and reproduces the fact that men are better matched for any given level of human capital.

RELATIONSHIP BETWEEN LABOR MARKET SORTING AND MARRIAGE MARKET SORTING. In Section 2, we documented a strong link between labor market and marriage market sorting in the data, whereby 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: Positive marriage sorting induces agents to exert the labor hours that ‘correspond’ to their human capital types, thereby ensuring a labor market match that fits their types. Note that consistent with our estimated model, here we proxy marriage market sorting in the data by spouses’ differences in *human capital s*-types (as opposed to the differences in education we used in Section 2). Similarly, labor market sorting is measured by the correlation of (*s*, *y*) (instead of (*x*, *y*)).

THE ROLE OF HOURS. A major feature of our model is that marriage and labor markets are linked in equilibrium, through the household’s time allocation choice. Here we show that the model

replicates salient features of the data, whereby hours are associated with both marriage and labor market outcomes. Figure 6, right panel, shows that in both the data (dashed) and the 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 $s_f \approx s_m$, at the vertical line). This is a natural prediction of our model: Spouses of similar human capital can better act on the hours complementarity in home production and align their hours more closely, 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 and wages. Figure 5 (right) shows the labor market matching function when *controlling* for hours worked. The difference in sorting across gender nearly vanishes in both the model (solid) and the data (dashed), relative to the left panel. Moreover, we find that increasing weekly labor hours from 30 to 40 raises the hourly wage by 5.6% in the model, which closely matches the effect of 4.2% in the data (see Section 2).

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)


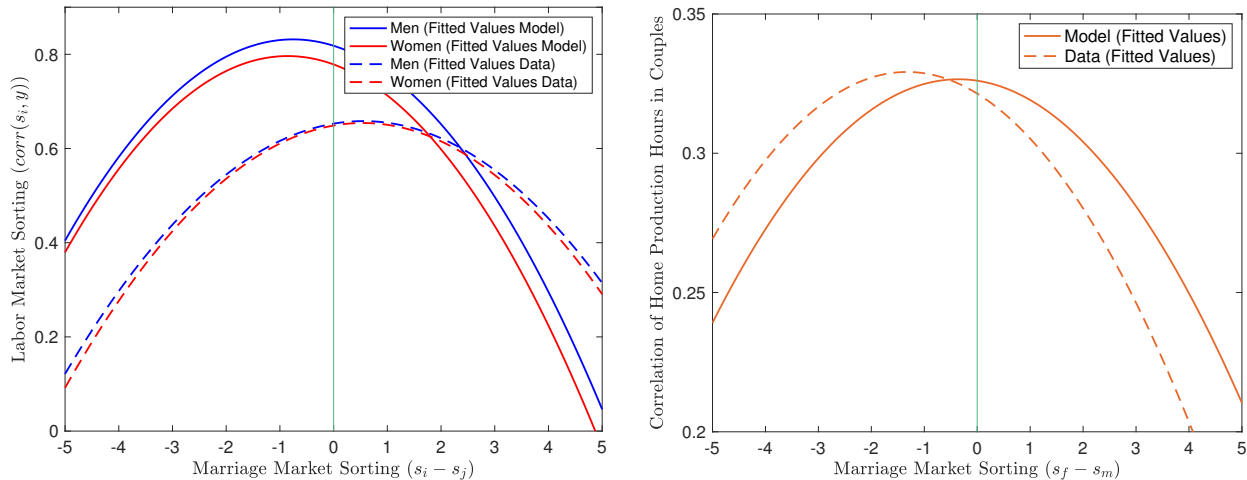


Figure 6: Labor Market Sorting and Marriage Market Sorting (left); Home Production Hours and Marriage Market Sorting (right)



5 Application: The Drivers of Inequality

In our main quantitative exercise, we use our model to shed new light on how home production complementarity affects gender disparities in the labor market and household income inequality. Our analysis focuses on two contexts: Germany in a recent cross-section (Section 5.1) and over time (Section 5.2).

5.1 Inequality in the Cross-Section

We first focus on a recent period, 2010-2016. We start by investigating the performance of our model in reproducing the observed inequality.

To assess the extent of inequality in data and model, we focus on four measures: the gender wage gap, the household wage variance, and its decomposition into between- and within-household components. These statistics are reported in Table 4.³⁶ While our model underestimates the level of the income variance (87 in the model versus 98 in the data), we capture the split of within- and between-household inequality quite well (47-53 split in the model vs. 50-50 in the data). Last, our model produces a sizable unconditional gender wage gap (24%), which slightly overestimates the observed gap (20%).

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

Table 4: Gender and Household Inequality

	Model	Data
Total Wage Variance	86.80	97.78
Within-household Wage Variance	40.47	49.18
... share in total variance	0.47	0.50
Between-household Wage Variance	46.33	48.60
... share in total variance	0.53	0.50
Gender Wage Gap	0.24	0.20

5.2 Inequality Changes over Time

Looking first at the data, inequality in Germany has changed significantly over the last few decades. In Figure 7, left panel, the turquoise bars show that the total 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. In the right panel, the turquoise bars show that positive marriage market

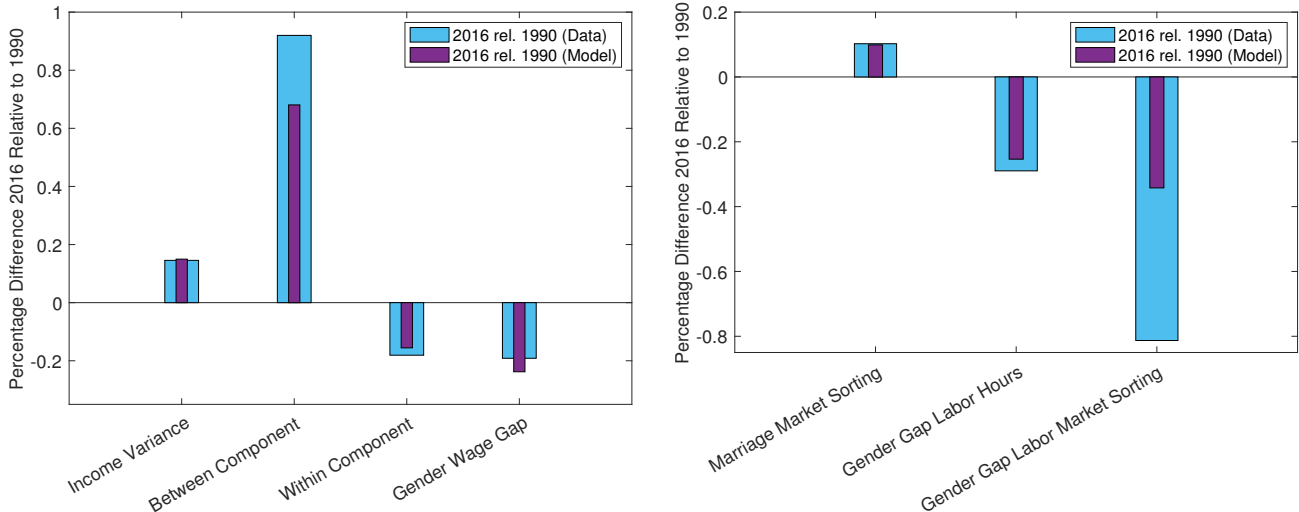
³⁶The gender wage gap is the difference in male and female mean wages over male mean wage. The within- and between component of wages are based on a standard variance decomposition of wages into variation between and within couples. Our measure of the gender wage gap includes all individuals, singles and in couples, conditional on employment. In turn, the total wage variance and its decomposition is based on the sample of couples, independent of employment status.

sorting increased by 10%, while the gender gap in labor hours fell by almost 30% and the gender gap in labor market sorting by around 80%.

We first discuss how our model explains these inequality shifts. Then we highlight the role changes in marriage and labor market sorting played in them.

THE ROLE OF MODEL PRIMITIVES. 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 then how these changes affected inequality.

Figure 7: Inequality Changes over Time (left); Sorting and Hours Changes over Time (right).



To this end, we re-estimate our model in an earlier period, 1990-1996, and compare it with our estimation for 2010-2016.³⁷ Table A.4 (Appendix E) shows that the model fits the targeted moments from 1990-1996 well. It also indicates that both the labor and marriage market underwent statistically significant changes over time (column 5). Regarding the untargeted inequality moments (Figure 7, left panel, turquoise bars), the model (purple bars) also does a good job in replicating over-time changes.

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 A.3, Appendix E. We observe significant changes in home production, with Germany in 2010-2016 being characterized by a lower ρ (dropping from 0.01 to -0.54 , which indicates increased complementarity in spouses' home hours) and a lower θ (dropping from 0.88 to 0.82, which implies that men became relatively more productive at home). Also, labor productivity wedge ψ has narrowed over time (increasing from 0.78 to 0.85, which reflects the rise in relative female productivity). This could be due to reduced discrimination or increased demand for skills in which women have a comparative advantage. These changes indicate that Germany has become an economy with more gender equality at both home and work. In turn, labor market technology has become more convex in effective human capital, resembling (effective) skill-biased technological change,

³⁷For 1990-1996, we reassess the skill and job distributions and re-estimate all parameters except for the scale of the labor supply shock, which we set to the level of 2010-2016 (Section 4.4.2). This is to avoid giving changes in shock distributions too prominent a role. We did have to free up the scale of the marriage shock to give the model a chance to match the data.

and it has a higher TFP than before.

How much of the documented changes in inequality can be explained by these changes in model parameters? Figure 8 provides a detailed decomposition. The purple bars again display the overall change in inequality produced by the model and account for *all* parameter changes over time. The remaining bars show 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).


The documented changes in home production technology significantly reduced gender disparities (gender wage gap and within-household inequality) as well as overall household inequality, while fueling between-household inequality. Figure 8 (orange bars) shows that home production changes were indeed the biggest driver of the decline in gender inequality. 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 13% (accounting for 54% of its drop). In turn, home production shifts put upward pressure on between-household inequality, accounting for almost 16% of its 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 A.1, Appendix E) reveals that changes in these parameters played an equally important role in these inequality shifts.

Regarding the documented labor market shifts, the effects of changes in the labor market wedge ψ on inequality (yellow bars, Figure 8)—while qualitatively similar to those of home production technology—were quantitatively smaller. Finally, changes in labor market technology (blue bars), and especially increases in γ_1 and A_z , fueled inequality across the board, pushing up both the between and within components of the household income variance and preventing gender inequality from falling even further. Thus, technological change in home production and in the labor market have moved inequality, and especially gender disparities, in opposite directions. The reason is that (effective) skill-bias technical change mainly benefited men, who work more in the labor market and are better matched to start with.

The comparative statics of our model, presented in Online Appendix OF, clarify the mechanism whereby the estimated changes in home production technology (particularly ρ and θ) and the labor wedge (ψ) increase gender equality.³⁸ While an increase in ψ boosts relative female labor productivity *directly*, there are also important *indirect* effects that stem from not only ψ but also θ and ρ and work through adjustments in hours. Changes in home production complementarities ρ and relative female home productivity θ as well as in female labor productivity ψ all induced women to work more, with the result

³⁸In theory, we cannot rule out multiple equilibria when considering the full problem. This potentially causes an issue for comparative statics whereby the effect of parameter changes on outcomes depends on equilibrium selection. However, the *analytical* properties of our model (i.e., in *each stage* of the model, uniqueness of equilibrium as well as intuitive comparative statics, which hold within the entire class of regular—and thus stable—equilibria) and the *numerical* properties (i.e., fast and monotone convergence to the equilibrium wage function for different classes of initial wage guesses, which suggests that we found a unique equilibrium, see Figure O.6, Online Appendix OC.5) indicate that multiplicity is not a major concern.

Figure 8: Mechanism behind Inequality Changes



that couples aligned their hours more closely. More aligned hours within couples led to a decline in the gender gap of labor hours, which in turn caused women to improve their effective human capital and thus sort relatively better in the labor market (thereby reducing the gender gap in labor sorting). Stronger complementarity in hours also bolsters the desire for positive sorting in marriage, which reinforces the shift toward more equal hours across gender. As a result, the gender wage gap and within-household income inequality drop even further, at the expense of larger between-household inequality.

In our model, a pattern arises whereby reductions in the gender wage gap and within-household income inequality (purple bars of Figure 7, left panel) go hand in hand with a decline in gender gaps in labor hours and labor market sorting but an increase in marriage market sorting (purple bars in the right panel). These inequality shifts and the underlying mechanism are consistent with the data (turquoise bars in the left and right panels of Figure 7).

Our evidence and estimates suggest that Germany underwent significant changes over the last decades toward an equilibrium that resembles the monotone one from our theory. This equilibrium is characterized by stronger home production complementarities and, consequently, increased marriage sorting as well as stronger comovements of spouses' hours, labor market sorting, and wages.

THE ROLE OF SORTING. Between 1990-1996 and 2010-2016, Germany saw large increases in positive sorting in the marriage market (by around 10%) and in the labor market (by 8%). We highlight the importance of accounting for marriage and labor market sorting in the model by quantifying how these changes in sorting affect inequality.

To do so, we compute the elasticity of each inequality outcome (gender wage gap, household income variance, within/between component) with respect to sorting in each market as

$$\left(\frac{\text{Ineq}(\hat{\Lambda}_t, \text{Sorting}_t) - \text{Ineq}(\hat{\Lambda}_t, \text{Sorting}_{t-1})}{\text{Ineq}(\hat{\Lambda}_t, \text{Sorting}_{t-1})} \right) / \left(\frac{\text{Sorting}_t - \text{Sorting}_{t-1}}{\text{Sorting}_{t-1}} \right),$$

where $t=2010-2016$ and $t-1=1990-1996$, $\hat{\Lambda}_t$ is the vector of estimated parameters in t , and $Sorting_t \in \{Labor\ Sorting_t, Marriage\ Sorting_t\}$ in t . The numerator gives the %-change in an inequality outcome between our baseline model 2010-2016, $Ineq(\hat{\Lambda}_t, Sorting_t)$, and that measure under a counterfactual model, $Ineq(\hat{\Lambda}_t, Sorting_{t-1})$, in which we apply the estimated parameters from baseline period t ($\hat{\Lambda}_t$) but keep sorting fixed at $t-1$ ($Sorting_{t-1}$). To implement the counterfactual that keeps marriage sorting fixed, we adjust σ_β to match the correlation across partners' types from the estimated model in $t-1$. In turn, to implement the counterfactual that keeps labor sorting fixed, we impose the labor market matching $\mu(\tilde{s})$ from the estimated model in $t-1$. The denominator is then computed as the %-change in sorting between period t and $t-1$, calculated as changes in the correlations of (s_m, s_f) for marriage sorting and of (s, y) for labor market sorting. The comparison between baseline and counterfactual models allows us to isolate the role of changes in sorting for inequality shifts.

Table 5 reports the elasticities. We find that *both* marriage and labor sorting have had mitigating impacts on gender inequalities (wage gap and within-household inequality) and have amplified between-household inequality. For instance, a 1% increase in marriage sorting has decreased within-household inequality by 0.12%, while it increased between-household inequality by 0.13%. 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 ultimately pay—across genders.

The effects of changes in labor sorting on inequality are even larger. A 1% rise in labor market sorting increased the between-household income variance by 0.44%. In turn, a 1% increase in labor sorting reduced the gender wage gap by 1.11% and within-household inequality by 0.77%. 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 (i.e., the gender gap in labor sorting has shrunk over time; Figure 7, right panel), which helped them catch up with men's pay. Stronger positive sorting between workers and jobs—when over-proportionally benefiting women—can spur gender convergence in labor market outcomes.

Table 5: Elasticity of Inequality with Respect to Sorting

	Income Variance	Between-Household Variance	Within-Household Variance	Gender Wage Gap
Marriage Market Sorting	0.01	0.13	-0.12	-0.01
Labor Market Sorting	-0.15	0.44	-0.77	-1.11

THE ROLE OF THE IMPACT OF LABOR MARKET SORTING ON THE MARRIAGE MARKET. In contrast to our analysis, a common approach in this literature is to keep the labor market in partial equilibrium and thereby neglect worker-job sorting. However, we show that there is a substantial quantitative impact of labor market sorting on the marriage market and, through this channel, on inequality. To do so, we proceed in two steps. We first compute the elasticity of marriage market

outcomes—marriage market sorting and households’ time allocation—with respect to labor market sorting. We find that a 1% increase in labor market sorting leads to a 0.79% increase in marriage market sorting and a 0.22% increase in the spousal correlation of home production hours. This suggests that labor market sorting plays a substantial quantitative role in shaping marriage market outcomes. Moreover, qualitatively, these patterns are in line with both our reduced-form evidence (Section 2) and the monotone equilibrium of our model (Section 3.4 and Online Appendix OB.3), whereby labor market sorting, marriage market sorting, and the complementarities of hours all move in the same direction.

In a second step, to gauge the quantitative importance of this channel for inequality, we compute the *counterfactual* elasticities of our four inequality measures with respect to labor market sorting while keeping marriage market outcomes *fixed* at the baseline level.³⁹ Our results show that *not* allowing for responses in marriage market outcomes dampens the impact of labor market sorting on inequality. In absolute terms, the elasticities of the between- and within-household wage variance with respect to labor market sorting are 30% and 14% *lower* in this counterfactual compared with the baseline elasticities (row 2 of Table 5). Thus, sorting in the labor market not only has a direct effect on inequality, but this impact is amplified by the endogenous marriage market response. A model with the labor market in partial equilibrium that lacks worker-job sorting would miss these interesting features.

6 The Sources of Home Production Complementarities

Given the prominence of home production complementarities in our analysis—driven mainly by ρ in the home production function—we provide additional evidence on the factors that drive them. We focus on two sources of complementarities and their changes over time: differences in complementarities within various tasks of home production and differences in complementarities across couples that differ in education. Our main conclusion will be that aggregate home production complementarities are driven by the first source—and, particularly, increasingly strong complementarities across spouses in childcare—while the second source has little bearing.

6.1 Heterogeneous Home Production Complementarities by Tasks

To zoom into the various home production tasks in more detail, we consult the German Time Use Survey (GTUS, 2023), a diary-based survey that features the most detailed time use data in Germany. We use the two waves of data, 1991/92 and 2012/13, that coincide with the time windows we study in the GSOEP. In Online Appendix OD.1, we describe the GTUS. Online Appendix OH shows that important time use facts are similar in the GTUS and GSOEP.

To assess the sources of aggregate home production complementarities empirically, we again rely on the correlation of home hours between spouses—the moment that informs home production complementarity ρ in estimation. Disaggregating home production into nine detailed tasks, we find that *childcare* is

³⁹These counterfactual elasticities of inequality with respect to labor market sorting are implemented as in row 2 of Table 5, except that we adjust σ_β to keep marriage market sorting (and hours correlation) at the baseline level (2010-2016).

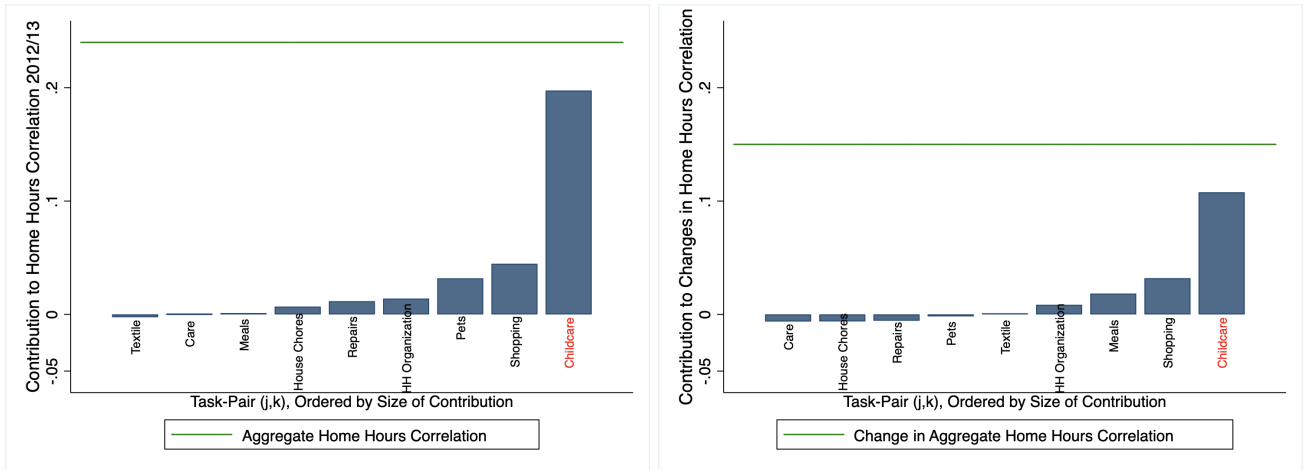
not only the largest home production category in 2012/13 (Figure O.11, right, Online Appendix OG.1) but it is also the most complementary activity by a wide margin: Within childcare, the correlation between spouses' hours is above 0.6 in 2012/13, compared with an aggregate correlation of 0.25 (Figure O.11, left). Also, over time, childcare is one of the home production tasks that experienced the largest increase in the hours correlation (Figure O.11, left).

To formalize this analysis, we perform a statistical decomposition of the aggregate correlation in home hours into the weighted sum of correlations of nine detailed home production tasks:

$$\text{Corr}(\ell_f, \ell_m) = \sum_{j=1}^9 \sum_{k=1}^9 \text{Corr}(\ell_{fj}, \ell_{mk}) \underbrace{\frac{\sqrt{\text{Var}(\ell_{fj}) \text{Var}(\ell_{mk})}}{\sqrt{\text{Var}(\ell_f) \text{Var}(\ell_m)}}}_{:=V_{jk}} = \sum_{j=1}^9 \sum_{k=1}^9 \text{Corr}(\ell_{fj}, \ell_{mk}) V_{jk}, \quad (8)$$

where V_{jk} is the 'weight' of the correlation between female home hours in task j (ℓ_{fj}) and male hours in task k (ℓ_{mk}). Figure 9 (left) shows that with a contribution of 85%, the (weighted) correlation of spouses' childcare is the main force behind the aggregate home production correlation today, consistent with the descriptive analysis above. We can further decompose over-time *changes* of the aggregate hours correlation (using formula (O.13) in Online Appendix OG.2): Around 73% of the total increase in spouses' home production correlation can be attributed to changes within childcare (weight and correlation), see Figure 9 (right).

Figure 9: Contribution of Detailed Home Production Tasks to Aggregate Home Production Correlation, Cross-section of 2012/13 (left) and over Time between 1991/92 and 2012/13 (right)



Notes: The left panel plots $\text{Corr}(\ell_{fj}, \ell_{mk}) V_{jk}$ for each task pair (j, k) s.t. $j = k$; see (8). The right panel plots over-time *changes* in $\text{Corr}(\ell_{fj}, \ell_{mk}) V_{jk}$, that is, $\overline{\text{Corr}(\ell_{fj}, \ell_{mk})} (V_{jk,t} - V_{jk,t-1}) + \bar{V}_{jk} (\text{Corr}(\ell_{fj}, \ell_{mk})_t - \text{Corr}(\ell_{fj}, \ell_{mk})_{t-1})$, for each task pair (j, k) , with $j = k$; see decomposition (O.13) in Online Appendix OG.2. In both panels we omit task pairs (j, k) s.t. $j \neq k$ from the figure since they do *not* significantly contribute to the aggregate correlation. Source: GTUS.

These results suggest that couples with children have stronger home production complementarities than those without children (also in line with the descriptive evidence of Figure O.12, left, Online Appendix OG.1) and also experienced the largest rise in complementarities. To investigate this system-

atically, we use decomposition (O.15) in Online Appendix OG.3 to break down the over-time change in the aggregate home production correlation into composition shifts (between-group changes) and within-group changes, where our ‘groups’ are couples *with* and *without* children. We find that changes in home hours correlation within couple-type—as opposed to composition shifts between couples with and without children—accounted for most of the change in the aggregate correlation (82%), driven by couples *with* children.

Finally, to complement this empirical analysis, we re-estimate our model allowing for heterogeneous home production parameters by couple-type—defined as couples with and without children. Details on this estimation exercise are in Online Appendix OG.4. Table O.17 displays the estimation results and Figure O.14 the model fit. Our main finding is that home production complementarities are significantly higher for couples with children, as indicated by a larger ρ in absolute value for these couples. Moreover, the gap between couples with and without children has increased over time.

Our results paint a consistent picture: Childcare is the most complementary home production task and this complementarity has increased over time. This shift, driven by couples with children, is a crucial driver behind our estimated increase in aggregate home production complementarities ($\rho \downarrow$) over time.

6.2 Heterogeneous Home Production Complementarities by Couples’ Education

An alternative hypothesis is that the aggregate home production complementarities conceal heterogeneity in complementarities across couples that differ in education. Our empirical evidence (Section 2) and our estimated model (Section 4.5) indicate that partners with similar education have a higher home production correlation compared with mixed couples (see also Figure O.12, right). This may suggest that home production complementarities are more pronounced for couples with similarly educated spouses. If complementarities have increased most strongly for that group or an increase in marriage market sorting puts more weight on those couples, then aggregate home production complementarities rise.

We first use decomposition (O.14) in Online Appendix OG.3 to split the aggregate correlation of spouses’ home production hours into within-/between-contributions from four couple-types (i.e., both partners are highly educated; both have low education; wife is highly educated but husband has low education, and vice versa). We find that the within-component explains 99.5% of the aggregate correlation, with the group of *same-educated* couples accounting for most of it (see Figure O.13, left). Moreover, decomposing changes in the aggregate correlation over time (based on (O.15)), we find that the overwhelming part (98%) of the rise in aggregate home production correlation can be accounted for by shifts within (as opposed to between) couple types. Zooming into these within-group changes, we find that couples with *similarly educated* partners drive most of this increase (Figure O.13, right).

A natural question then is whether home production complementarities ρ are heterogeneous across couple types in our model and whether a larger shift in ρ for positively assorted couples drives the aggregate change in home production complementarity over time. To address this question, we re-estimate our model to allow for heterogeneous home production parameters by couple education, where

we again distinguish between four couple types. Details on the estimation are in Online Appendix OG.4. We pin down four couple-type-specific home production complementarities, $\rho_{\mathcal{G}}$, by targeting the home production correlations across couple types \mathcal{G} . The estimation results are in Table O.18 and the model fit in Figure O.15.

Our main finding is that despite the fact that similarly educated couples have higher home production correlations than couples with mixed education, the estimates of $\rho_{\mathcal{G}}$ are *not* statistically different across couple types. Thus, our model does not require heterogeneity in home production complementarities to generate heterogeneous home production correlations. Instead, even with an economy-wide ρ , it can endogenously generate the higher home production correlation in couples with similarly educated spouses. Based on this analysis, we conclude that heterogeneity in home production across couples that differ in education is *not* crucial for understanding aggregate home production complementarities.

The analysis in this section suggests that the aggregate increase in home production complementarities from our baseline analysis is driven by changes in complementarities within home production tasks. Specifically, a boost to spousal complementarities in childcare—the task with the largest share within home production—appears to be crucial. In line with this finding, our new estimations reveal that couples with children have larger complementarities and changes therein over time. In contrast, we find no evidence for heterogeneity in home production complementarities across couples that differ in education or differential changes in couple-specific complementarities over time.

7 Conclusion

Employers value workers not only for their skills but also for how many hours they are willing to work. Therefore, labor supply decisions affect the jobs workers can get and thus labor market sorting. But if these labor supply choices are made within the household and depend on the characteristics of *both* spouses, then marriage market sorting affects labor market sorting and, ultimately, wages and inequality.

At the center of this paper is the interplay between labor and marriage markets and how it shapes inequality across gender and within/between households. 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 the model, the nature of home production—whether partners’ hours are complements or substitutes—shapes marriage market sorting, labor supply choices, and labor market sorting.

We then examine the nature of home production in the data. To this end, we estimate our model on data from modern Germany and find that spouses’ home hours are strategic *complements*. This complementarity reinforces positive sorting in both markets and the comovement of spouses’ labor hours. This is in contrast to what would happen in a ‘traditional’ economy based on substitution in home production and the specialization of spouses. By investigating the critical drivers behind inequality, 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 between partners. Home production complementarities induce spouses to make similar time allocation choices regarding work in the market and at home; they also increase marriage sorting and reduce the gender gap in labor sorting. Both channels further mitigate gender disparities.

Our main quantitative exercise analyzes how our model rationalizes changes in inequality over time. We find that over recent decades, the home production hours of spouses have become more complementary. 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. To highlight the unique feature of our model, we show that sorting on both markets has significant quantitative effects on inequality: *Both* stronger marriage market sorting and labor market sorting over time have amplified overall inequality and between-household inequality. But they have also had a mitigating impact on gender inequalities (wage gap and within-household inequality), which reveals a new role of sorting for gender convergence in pay.

Given the prominence of home production complementarities in our analysis, we further investigate their main sources. We find that a boost to spousal complementarities in *childcare*—the task with the largest share within home production—is the crucial driver behind the estimated aggregate increase in home production complementarities over time.

This paper opens a new research agenda on the interplay between labor and marriage sorting in an equilibrium setting and its implications for gender gaps and inequality. We expect that extensions of our model can be used to study how sorting in both markets impacts households' responses to labor market shocks and risk-sharing arrangements. We plan to address these issues in future work.

Data Availability Statement. The data and code underlying this research is available on Zenodo at <https://dx.doi.org/10.5281/zenodo.10137965>.

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.
- M. Aguiar and E. Hurst. Life-cycle prices and production. *American Economic Review*, 97(5):1533–1559, 2007.
- 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.
- G. S. Becker. Human capital, effort, and the sexual division of labor. *Journal of Labor Economics*, 3(1):S33–S58, 1985.
- 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. *The Quarterly Journal of Economics*, 137(3):1901–1962, 2022.
- 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, 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.
- V. Chernozhukov, A. Galichon, M. Henry, and B. Pass. Identification of hedonic equilibrium and nonseparable simultaneous equations. *Journal of Political Economy*, 129(3):842–870, 2021.
- 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, 2017.
- 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. NBER Working Paper 26932, 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.
- M. Costa Dias, R. Joyce, and F. Parodi. The gender pay gap in the UK: children and experience in work. *Oxford Review of Economic Policy*, 36(4):855–881, 2021.
- G. Cubas, C. Juhn, and P. Silos. Coordinated Work Schedules and the Gender Wage Gap. *The Economic Journal*, 133(651):1036–1066, 11 2022.
- A. S. Edlin and C. Shannon. Strict monotonicity in comparative statics. *Journal of Economic Theory*, 81(1): 201–219, 1998.
- 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. *American Economic Journal: Macroeconomics*, 14(3):543–90, 2022.
- R. Fernández and R. Rogerson. Sorting and long-run inequality. *The Quarterly Journal of Economics*, 116(4):1305–1341, 2001.
- R. Fernández and J. C. Wong. Divorce risk, wages and working wives: A quantitative life-cycle analysis of female labour force participation. *The Economic Journal*, 124(576):319–358, 2014.
- R. Fernández and J. C. Wong. Free to leave? a welfare analysis of divorce regimes. *American Economic Journal: Macroeconomics*, 9(3):72–115, July 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.
- German Time Use Survey (GTUS). Data for 1991/92 and 2012/13. German Federal Statistical Office. doi:10.21242/63911.1992.00.00.4.1.0 and 10.21242/63911.2013.00.00.4.1.0, 2023.
- C. Goldin. A grand gender convergence: Its last chapter. *American Economic Review*, 104(4):1091–1119, 2014.
- C. Goldin. *Career and Family: Women’s Century-Long Journey toward Equity*. Princeton University Press, 2021.
- 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, 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.
- A. Hall, A. Seifer, and M. Tiemann. Bibb/baua employment survey of the working population on qualification and working conditions in germany 2012. version 6.0. doi:10.7803/501.12.1.1.60, Research Data Center at BIBB (ed.); GESIS Cologne (data access); Bonn: Federal Institute for Vocational Education and Training., 2020.
- 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. S¸ogaard. Children and gender inequality: Evidence from Denmark. *American Economic Journal: Applied Economics*, 11(4), 2019.
- 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. *Journal of Political Economy*, (forthcoming), 2023.
- J. Lise and F. Postel-Vinay. Multidimensional skills, sorting, and human capital accumulation. *American Economic Review*, 110(8):2328–76, 2020.
- J. Lise and S. Seitz. Consumption inequality and intra-household allocations. *The Review of Economic Studies*, 78(1):328–355, 01 2011.
- 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.
- P. Milgrom and I. Segal. Envelope theorems for arbitrary choice sets. *Econometrica*, 70(2):583–601, 2002.
- P. Milgrom and C. Shannon. Monotone comparative statics. *Econometrica*, 62(1):157–80, 1994.
- E. Moschini. Child care subsidies and child skill accumulation in one- and two-parent families. *American Economic Journal: Macroeconomics*, Forthcoming.
- A. Pakes and D. Pollard. Simulation and the asymptotics of optimization estimators. *Econometrica*, 57(5): 1027–1057, 1989.
- L. Pilossoph and S. L. Wee. Household search and the marital wage premium. *American Economic Journal: Macroeconomics*, 13(4):55–109, 2021.
- L. Pilossoph and S. L. Wee. Assortative matching and household income inequality: A structural approach. 2023.
- A. Reynoso. The impact of divorce laws on the equilibrium in the marriage market. 2022.
- D. Rohrbach-Schmidt and A. Hall. Bibb/baua-erwerbstaetigenbefragung 2012. BIBB-FDZ Daten- und Methodenberichte 1/2013 Version 6.0., Bonn: BIBB. ISSN 2190-300X, 2020.
- S. Rosen. Hedonic prices and implicit markets: Product differentiation in pure competition. *Journal of Political Economy*, 82(1):34–55, 1974.
- P. Rupert, R. Rogerson, and R. Wright. Estimating substitution elasticities in household production models. *Economic Theory*, 6(1):179–193, 1995.
- Socioeconomic Panel (SOEP) Data for Years 1984-2018, Version 35, soep, doi:10.5684/soep.v35., 2019.
- 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: Proof of Proposition 1

We focus on any *regular equilibrium*, in which household problem (5) has an interior and continuous solution that is unique, and \tilde{N} is atomless (for the formal definition and details, see Online Appendix OB.2).

LABOR MARKET PROPERTIES. Supermodularity of z implies positive sorting in the labor market. Moreover, in a regular equilibrium, we have the following properties of \tilde{N} : First, it is defined over an interval. This follows from $x \in [0, \bar{x}]$, together with the regularity assumptions that the solution to the household problem is interior and the functions $h_i, i \in \{f, m\}$ are continuous on $[0, \bar{x}] \times [0, \bar{x}]$, and so the solution to the household problem satisfies $h_i \in [\underline{h}_i, \bar{h}_i], i \in \{f, m\}$ with $\underline{h}_i > 0, \bar{h}_i < 1$. As a consequence, the range of the effective type function, $e(\cdot, \cdot)$, is an interval. Second, \tilde{N} is atomless.

Then, under positive sorting, $\mu(\tilde{x}) = G^{-1}(N(\tilde{x}))$, with $\mu' > 0$. Moreover, the wage function w is given by $w(\tilde{x}) = \int_0^{\tilde{x}} z_{\tilde{x}}(t, \mu(t)) dt$, which is strictly increasing and strictly convex, where strict convexity follows since z is strictly supermodular in (\tilde{x}, y) and convex in \tilde{x} for each y .⁴⁰

MARRIAGE MARKET PROPERTIES. Turning to the marriage stage, consider any couple (x_m, x_f) who jointly chooses hours h_m and h_f . The household's problem is

$$\Phi(x_m, x_f) = \max_{h_m, h_f \in [0, 1]} (w(e(x_m, h_m)) + w(e(x_f, h_f)) + 2p(1 - h_m, 1 - h_f)),$$

where we replaced $\tilde{x}_m = e(x_m, h_m)$ and $\tilde{x}_f = e(x_f, h_f)$, and where we denote the value of this problem by $\Phi(x_m, x_f)$. If hours are strictly increasing in (x_m, x_f) , then $\Phi(x_m, x_f)$ is strictly supermodular and PAM emerges in the marriage market. We turn to the strict monotonicity of the hours functions next.

PROPERTIES OF THE HOURS FUNCTIONS. We will show that h_m strictly increases in x_m and x_f . To do so, write the household's problem as follows:

$$\max_{h_m \in [0, 1]} \left(w(e(x_m, h_m)) + \max_{h_f \in [0, 1]} (w(e(x_f, h_f)) + 2p(1 - h_m, 1 - h_f)) \right).$$

That is, we split the joint maximization w.r.t. (h_m, h_f) into two maximization problems.

Let V be the value of the inner maximization problem, that is

$$V(x_f, h_m) = \max_{h_f \in [0, 1]} (w(e(x_f, h_f)) + 2p(1 - h_m, 1 - h_f)).$$

⁴⁰As is known in models with pre-match investment (see, e.g., Cole et al., 2001), a potential issue arises if off-equilibrium hours choices are not priced. In our monotone equilibrium, however, this is something we can address in a straightforward way. First note that, as argued above, effective types in equilibrium are distributed on an interval, $\tilde{x} \in [0, e(\bar{x}, \bar{h})]$, and the wage function $w(\tilde{x}) = \int_0^{\tilde{x}} z_{\tilde{x}}(t, \mu(t)) dt$ is defined on that interval. This also implies that all off-equilibrium hours choices $0 \leq h < \underline{h}$ are priced by our wage function. Finally, in order to price the off-equilibrium choices $\bar{h} < h \leq 1$, we follow Cole et al. (2001) and 'extend' the wage function using positive sorting between (\tilde{x}, y) . To do so, note that firms' payoff is given by $\pi(y) = \int_0^y z_y(\mu^{-1}(t), t) dt$, independently of whether the effective type $\mu^{-1}(y)$ is on or off equilibrium. But then, a worker with skill x who chooses off-equilibrium hours $h > \bar{h}$ and whose effective type possibly satisfies $\tilde{x} > e(\bar{x}, \bar{h})$ has the well-defined payoff $w(\tilde{x}) = z(\tilde{x}, \mu(\tilde{x})) - \pi(\mu(\tilde{x}))$. At the core of this argument is that positive sorting between workers and firms holds—and thus assignment $\mu(\tilde{x})$ is well-defined—even when considering off-equilibrium effective types.

We will first focus on the outer maximization problem

$$\max_{h_m \in [0,1]} \left(w(e(x_m, h_m)) + V(x_f, h_m) \right),$$

taking as given the function V . We will provide conditions on V under which this problem satisfies the strict single crossing property (SSCP) in $(h_m, (x_m, x_f))$. Under SSCP, all the selections from the optimal correspondence are increasing in x_m and x_f (Milgrom and Shannon, 1994, Theorem 4').⁴¹ And since in a regular equilibrium the hours functions satisfying the households' optimality conditions are unique, this unique solution is increasing in x_m and x_f as well. Then, we will show that h_m is actually *strictly* increasing in both attributes.

Since the objective function is additively separable in (x_m, h_m) and (x_f, h_m) , it follows that SSCP holds if each term is strictly supermodular (i.e., satisfies strictly increasing differences). Given the properties of e and w , we have that $w(e(\cdot, \cdot))$ is strictly supermodular in (x_m, h_m) since it is the composition of a convex function with a strictly supermodular one. So if V is also strictly supermodular in (x_f, h_m) , it will follow that h_m is increasing in both x_m and x_f (we will verify the supermodularity of V below).

To show that h_m is *strictly* increasing, we will use Edlin and Shannon (1998), Theorem 1 and Corollary 1. Since in any regular equilibrium, the optimal h_m is interior for all (x_m, x_f) (which in our case materializes due to the Inada conditions on p), the first-order condition of the household problem characterizes the optimal choices. This first-order condition is given by:

$$w'(e(x_m, h_m))e_{h_m}(x_m, h_m) + V_{h_m}(x_f, h_m) = 0.$$

Consider $\hat{x}_m > x_m$. Then since $w(e(\cdot, \cdot))$ is strictly supermodular in (x_m, h_m) , it follows that

$$w'(e(\hat{x}_m, h_m))e_{h_m}(\hat{x}_m, h_m) + V_{h_m}(x_f, h_m) > 0,$$

so the optimal hours for \hat{x}_m , say \hat{h}_m , are different than h_m (which are the optimal hours for x_m). But since we know from SSCP that $\hat{h}_m \geq h_m$, it follows that we must have $\hat{h}_m > h_m$. Hence, h_m is strictly increasing in x_m . The same holds for x_f under V strictly supermodular. Thus, h_m is *strictly increasing* in both (x_m, x_f) .

Let us now consider the inner maximization problem:

$$V(x_f, h_m) = \max_{h_f \in [0,1]} (w(e(x_f, h_f)) + 2p(1 - h_m, 1 - h_f)).$$

We first obtain that all selections of the correspondence of maximizers are increasing, as above. Note that both terms are strictly supermodular: the first one in (x_f, h_f) (since $w(e(\cdot, \cdot))$ is the composition of a supermodular and a convex function) and the second in (h_m, h_f) (under our assumption that p is

⁴¹SSCP holds if for all $(x'_m, x'_f) > (x''_m, x''_f)$ (in the standard vector order) and $h'_m > h''_m$, $w(e(x'_m, h'_m)) + V(x'_f, h'_m) - (w(e(x''_m, h''_m)) + V(x'_f, h''_m)) \geq 0$ implies $w(e(x'_m, h'_m)) + V(x'_f, h'_m) - (w(e(x'_m, h''_m)) + V(x'_f, h''_m)) > 0$.

strictly supermodular in (ℓ_m, ℓ_f) . So as above, SSCP in $(h_f; h_m, x_f)$ holds and thus the unique solution (in any regular equilibrium) to the household problem, h_f , is increasing in both x_f and h_m . And since h_m is strictly increasing in x_m (see above), we obtain that h_f is increasing in x_m as well (by taking the composition of functions).

We now show, as above, that given that the solution of the inner maximization problem is interior for all (x_f, h_m) —again ensured by the Inada conditions on p —it must be *strictly* increasing. The FOC is

$$w'(e(x_f, h_f))e_{h_f}(x_f, h_f) - 2p\ell_f(1 - h_m, 1 - h_f) = 0.$$

Consider $\hat{x}_f > x_f$. Since $w(e(\cdot, \cdot))$ is strictly supermodular, we have that

$$w'(e(\hat{x}_f, h_f))e_{h_f}(\hat{x}_f, h_f) - 2p\ell_f(1 - h_m, 1 - h_f) > 0.$$

so the optimal hours of \hat{x}_f , say \hat{h}_f , are different than h_f (which are the optimal hours of x_f). But since we know from SSCP that $\hat{h}_f \geq h_f$, it follows that $\hat{h}_f > h_f$. Hence, h_f is strictly increasing in x_f . Similarly, h_f is strictly increasing in h_m since p is strictly supermodular. Thus, h_f is *strictly increasing* in both x_f and h_m (and, due to the strict monotonicity of h_m in x_m , it is also strictly increasing in x_m).

To complete the proof, it remains to show that V is strictly supermodular in (x_f, h_m) and differentiable in h_m . Note that

$$V(x_f, h_m) = w(e(x_f, h_f(x_f, h_m))) + 2p(1 - h_m, 1 - h_f(x_f, h_m)).$$

By the Envelope Theorem (note that we satisfy the assumptions of [Milgrom and Segal, 2002](#), Corollary 4(iii)), especially since the solution of the household problem is unique in a regular equilibrium), V is differentiable in h_m with $V_{h_m}(x_f, h_m) = -2p\ell_m(1 - h_m, 1 - h_f(x_f, h_m))$. Since p is strictly supermodular and h_f strictly increasing in x_f , it follows that V is strictly supermodular. Hence, the premise made above, when analyzing the choice of h_m in the outer maximization problem, holds.

Since we have shown that the hours functions are strictly increasing in each attribute for *any* couple type, they are also strictly increasing along the equilibrium marriage market assignment $(\eta(x_f), x_f)$.

Finally, since hours are strictly increasing in the attributes, it follows that \tilde{N} is atomless, thus justifying the premise made in the labor market stage, which completes the proof. \square

B Quantitative Extension of the Model

In the quantitative extension of our model described in Section 3.5, we augment the baseline model to allow for shocks and unobserved skills. We here display the key decisions in more detail.

LABOR MARKET. A firm of type y now chooses a worker with human capital \tilde{s} (instead of \tilde{x}):

$$\max_{\tilde{s}} z(\tilde{s}, y) - w(\tilde{s}).$$

HOUSEHOLD PROBLEM. In any given couple, partners now maximize utility *plus* labor supply shock:

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

where we denote by p^M the home production technology of couples (M stands for married).

Similarly, the consumption-time allocation problem of singles $i \in \{f, m\}$ is given by

$$\begin{aligned} \max_{c_i, h_i} \quad & u(c_i, p^U(1 - h_i)) + \delta^{h_i} \\ \text{s.t.} \quad & c_i - w(\tilde{s}_i) = 0, \end{aligned} \tag{A.2}$$

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

MARRIAGE MARKET. The marriage problem of a man with human capital s_m now reads

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

where the choice of marrying a woman of any human capital type $s = s_f \in \mathcal{S}$ must be weighed against the choice of remaining single $s = \emptyset$ (i.e., $\Phi(s_m, \emptyset, v(s_m, \emptyset))$ 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 (details are in Online Appendix OC). Since marriage market matching is no longer pure (due to both the discreteness of 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). We then denote by $\eta(s_m, s_f)$ the fraction of couples with human capital types (s_m, s_f) , where $\sum_{(s_m, s_f) \in \{\mathcal{S} \cup \emptyset\}^2} \eta(s_m, s_f) = 1$.

C Identification

C.1 Identification of the Worker and Job Distributions

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 effects in a panel wage regression. See Section 4.4 and Online Appendix OE.3 and OE.4 for the details on estimation.

C.2 Identification of the Model Parameters: Proof of Proposition 2

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 rely on [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 non-linearities of our matching model, which are an important source of identification just as in [Ekeland et al. \(2004\)](#). Recall the firm’s optimality condition satisfies:

$$w'(\tilde{s}) = z_{\tilde{s}}(\tilde{s}, \mu(\tilde{s})). \quad (\text{A.3})$$

This equation can be used to identify the parameters of interest. We treat $w'(\tilde{s})$ as observed (it can be obtained as the derivative of the kernel regression of w (observed) on \tilde{s} in the subsample of men, where $\psi = 1$ by assumption), and denote its estimate by $\hat{w}_{\tilde{s}}$.

We can then identify the production function from FOC [\(A.3\)](#) after applying a log transformation and taking into account measurement error:

$$\log(\hat{w}_{\tilde{s}}(\tilde{s})) = \log(z_{\tilde{s}}(\tilde{s}, \mu(\tilde{s}))) + \epsilon \quad (\text{A.4})$$

where, for concreteness, we assume the functional form $z(\tilde{s}, y) = A_z \tilde{s}^{\gamma_1} y^{\gamma_2} + K$ (see [Section 4.1](#)), and where we treat \tilde{s} and the equilibrium 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, which has mean zero and is uncorrelated with the right-hand-side (RHS) variables. Regression [\(A.4\)](#) 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 $\underline{\tilde{s}}$, who—due to PAM in the labor market—matches with the lowest firm type $y = 0$ and thus $z(\underline{\tilde{s}}, 0) = K$, meaning any positive wage $w(\underline{\tilde{s}}) > 0$ can only be attributed to K , i.e., $w(\underline{\tilde{s}}) = K$.

Identification of the Female Productivity Wedge. We can identify ψ from the within (s, h) -type gender wage gap (i.e., from men and women with the same s and h). Denote the gender wage gap within the group of individuals with $(s, h) = (\hat{s}, \hat{h})$ 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 for (\hat{s}, \hat{h}) can be expressed as:

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

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 the matching between workers of any \tilde{s} with firms of type y .

Thus, we can consider labor market 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:

$$w(\psi\hat{s}\hat{h}) = w(\hat{s}\hat{h})(1 - gap(\hat{s}, \hat{h})) \quad (\text{A.5})$$

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 of ψ with $w(\psi\hat{s}\hat{h}) = K$ for $\psi = 0$ and $w(\psi\hat{s}\hat{h}) = w(\hat{s}\hat{h})$ for $\psi = 1$.

Hence, one of the following is true: if there is an interior gap, $0 < gap(\hat{s}, \hat{h}) < 1 - K/w(\hat{s}\hat{h})$, then by the Intermediate Value Theorem, there exists a unique $\psi \in (0, 1)$ for which (A.5) holds; or, if there is no gap, $gap(\hat{s}, \hat{h}) = 0$, then $\psi = 1$; finally, if there is a maximal gap, $gap(\hat{s}, \hat{h}) = 1 - K/w(\hat{s}\hat{h})$, then $\psi = 0$. Thus, ψ is identified from the gender wage gap of agents with the same (s, h) -bundle.

Identification of the Scale of the Labor Supply Shock. Recall that the choice set of singles differs from that of couples. We denote the hours alternatives that a decision maker $t \in \{M, U\}$ can choose by $\mathbf{h}^t \in \{\mathcal{H} \cup \emptyset\}^2 := \{\{0, \dots, 1\} \cup \emptyset\}^2$ where:

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

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

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

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

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

$$\pi_{\mathbf{s}}^t(\mathbf{h}^t) = \frac{\exp(\bar{u}_{\mathbf{s}}^t(\mathbf{h}^t)/\sigma_{\delta})}{\sum_{\tilde{\mathbf{h}}^t \in \{\mathcal{H} \cup \emptyset\}^2} \exp(\bar{u}_{\mathbf{s}}^t(\tilde{\mathbf{h}}^t)/\sigma_{\delta})},$$

which follows from our assumption on the labor supply shock distribution (type-I extreme value).

Let $\mathbf{h}^U = \mathbf{0} := (0, \emptyset)$ denote the hours of 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:

$$\begin{aligned} \frac{\pi_{\mathbf{s}}^U(\mathbf{h}^U)}{\pi_{\mathbf{s}}^U(\mathbf{0})} &= \frac{\exp(\bar{u}_{\mathbf{s}}^U(\mathbf{h}^U)/\sigma_{\delta})}{\exp(\bar{u}_{\mathbf{s}}^U(\mathbf{0})/\sigma_{\delta})} \\ \Leftrightarrow \log\left(\frac{\pi_{\mathbf{s}}^U(\mathbf{h}^U)}{\pi_{\mathbf{s}}^U(\mathbf{0})}\right) &= \frac{\bar{u}_{\mathbf{s}}^U(\mathbf{h}^U) - \bar{u}_{\mathbf{s}}^U(\mathbf{0})}{\sigma_{\delta}} = \frac{w(s_m h_m) + p^U(1 - h_m) - p^U(1)}{\sigma_{\delta}}, \end{aligned}$$

where the wage from not working is zero and where h_m are the hours associated with this single's 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 the difference in relative choices of these two single men is:

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

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 (i.e., the effect of men's human capital on wages given the hours choice $\mathbf{h}^U \neq \mathbf{0}$) is also observed. Thus, σ_δ is identified from the relative hours choices of single men.

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 choosing hours $\mathbf{h}^M \neq \mathbf{1}$ versus $\mathbf{h}^M = \mathbf{1}$ for married couple \mathbf{s} as:

$$\begin{aligned} \frac{\pi_{\mathbf{s}}^M(\mathbf{h}^M)}{\pi_{\mathbf{s}}^M(\mathbf{1})} &= \frac{\exp(\bar{u}_{\mathbf{s}}^M(\mathbf{h}^M)/\sigma_\delta)}{\exp(\bar{u}_{\mathbf{s}}^M(\mathbf{1})/\sigma_\delta)} \\ \log \left(\frac{\pi_{\mathbf{s}}^M(\mathbf{h}^M)}{\pi_{\mathbf{s}}^M(\mathbf{1})} \right) &= \frac{w(\psi s_f h_f) - w(\psi s_f) + w(s_m h_m) - w(s_m) + 2p^M(1 - h_m, 1 - h_f)}{\sigma_\delta}, \end{aligned} \quad (\text{A.6})$$

where we used that $2p^M(0, 0) = 0$ by assumption in our quantitative model. Note that the LHS of (A.6) is observed; on the RHS, wages of men and women with types (s_m, s_f) conditional on hours are also observed, and σ_δ is known at this stage. Thus, home production function p^M is non-parametrically identified since we can specify (A.6) for all hours alternatives $\mathbf{h}^M \neq \mathbf{1}$ chosen in the data. Note that we can identify p^M from a couple of any type $\mathbf{s} = (s_m, s_f)$. By a similar argument, p^U for singles is identified.

Identification of the Scale of the Marriage Taste Shock. We now show that σ_β is identified given that the parameters of the utilities are identified (as we have shown above).

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

$$\eta_{s_m}(s_f) = \frac{\exp(\Phi(s_m, s_f, v(s_m, s_f)))/\sigma_\beta}{\sum_{s'_f \in \{S \cup \emptyset\}} \exp(\Phi(s_m, s'_f, v(s_m, s'_f)))/\sigma_\beta},$$

where, as before, we denote by $\Phi(s_m, s_f, v(s_m, s_f))$ the expected value of a man with human capital s_m from being married to a woman of type s_f and paying her the transfer $v(s_m, s_f)$ (and where $\Phi(s_m, \emptyset, v(s_m, \emptyset))$ in the denominator denotes the expected value of being single with $v(s_m, \emptyset) = 0$). Value $\Phi(s_m, s_f, v(s_m, s_f))$ is given by

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

Consider a man with s_m . His ratio of choice probabilities regarding a woman s_f and staying single is:

$$\begin{aligned}
\log \left(\frac{\eta_{s_m}(s_f)}{\eta_{s_m}(\emptyset)} \right) &= \log \left(\frac{\eta(s_f, s_m)}{\eta(\emptyset, s_m)} \right) = \frac{\Phi(s_m, s_f, v(s_m, s_f)) - \Phi(s_m, \emptyset, v(s_m, \emptyset))}{\sigma_\beta} \\
&= \frac{\sigma_\delta}{\sigma_\beta} \left[\kappa + \log \left(\sum_{\mathbf{h}^M \in \mathcal{H}^2} \exp(\bar{u}_s^M(\mathbf{h}^M)/\sigma_\delta) \right) \right] - \frac{v(s_m, s_f)}{\sigma_\beta} \\
&\quad - \frac{\sigma_\delta}{\sigma_\beta} \left[\kappa + \log \left(\sum_{\mathbf{h}^U \in \mathcal{H}} \exp(\bar{u}_{s_m}^U(\mathbf{h}^U)/\sigma_\delta) \right) \right],
\end{aligned}$$

where $\eta_{s_m}(s_f)n_{s_m} = \eta(s_f, s_m)$ is the mass of couples of type (s_f, s_m) in equilibrium—implicitly assuming marriage market clearing—and $\eta_{s_m}(\emptyset)n_{s_m} = \eta(\emptyset, s_m)$ is the mass of single men of type s_m ; and n_{s_m} is the exogenous probability mass function for male human capital corresponding to cdf N_s .

Forming the analogue relative choice probability for a woman of type s_f yields:

$$\log \left(\frac{\eta_{s_f}(s_m)}{\eta_{s_f}(\emptyset)} \right) = \log \left(\frac{\eta(s_f, s_m)}{\eta(s_f, \emptyset)} \right) = \frac{v(s_m, s_f)}{\sigma_\beta} - \frac{\sigma_\delta}{\sigma_\beta} \left[\kappa + \log \left(\sum_{\mathbf{h}^U \in \mathcal{H}} \exp(\bar{u}_{s_f}^U(\mathbf{h}^U)/\sigma_\delta) \right) \right],$$

where $\eta_{s_f}(s_m)n_{s_f} = \eta(s_f, s_m)$ is the mass of couples of type (s_f, s_m) in equilibrium (and $\eta_{s_f}(\emptyset)n_{s_f} = \eta(s_f, \emptyset)$ is the mass of single women of type s_f); and n_{s_f} is the exogenous probability mass function for female human capital corresponding to cdf N_s .

We combine these relative choice probabilities of s_f and s_m by eliminating transfer $v(s_m, s_f)$:

$$\log \left(\frac{\eta(s_f, s_m)^2}{\eta(s_f, \emptyset)\eta(\emptyset, s_m)} \right) = \frac{1}{\sigma_\beta} \text{surplus}(s_m, s_f), \tag{A.7}$$

where we defined match surplus of couple type (s_f, s_m) as

$$\begin{aligned}
\text{surplus}(s_m, s_f) &:= \sigma_\delta \left[\kappa + \log \left(\sum_{\mathbf{h}^M \in \mathcal{H}^2} \exp(\bar{u}_s^M(\mathbf{h}^M)/\sigma_\delta) \right) \right. \\
&\quad \left. - \left(2\kappa + \log \left(\sum_{\mathbf{h}^U \in \mathcal{H}} \exp(\bar{u}_{s_m}^U(\mathbf{h}^U)/\sigma_\delta) \sum_{\mathbf{h}^U \in \mathcal{H}} \exp(\bar{u}_{s_f}^U(\mathbf{h}^U)/\sigma_\delta) \right) \right) \right].
\end{aligned}$$

The relative marriage frequencies on the LHS of (A.7) are observed. Moreover, all objects on the RHS are either observed (wages) or identified at this stage (home production function p and σ_δ), except σ_β . We can solve this equation for σ_β , giving a unique solution. Thus, σ_β is identified. \square

D Estimation

D.1 External Calibration

We calibrate the adjustment to the value of single (see footnote 29) and constant K outside of the main model estimation. Here we provide details and summarize the calibrated parameters in Table A.1.

Calibration of the Value of Singlehood. To match the fraction of singles in the data, we need more variation in the value of being single than what the marriage market preference shocks—which are driven by an extreme-value distribution with scale parameter σ_β —can provide while still matching the observed marriage market sorting. This could in principle be achieved with a nested-logit specification, which allows for different scale parameters for the ‘marriage nest’ and the ‘single nest’. However, because the single nest is degenerate, this approach would face the problem that the scale parameter of the degenerate nest cannot be uniquely identified (see Hunt, 2000).

To circumvent the nested logit specification while still guaranteeing enough variation in the value of being single to match the extensive margin of marriage, we calibrate an adjustment to the economic value of being single *outside* of the main estimation. This adjustment is a shock $\zeta \in \mathcal{Z}$, drawn by each individual from cdf Z , where—for convenience—we assume that Z is an extreme-value distribution with location zero and scale σ_ζ . The economic value of being single for an individual is then $w + p^U + \zeta$. By estimating only one scale parameter σ_β internally, this approach preserves the logit form of the model.

We calibrate σ_ζ in an iterative procedure outside of the internal estimation as follows.

- Step 0. Start with an internal estimation of the model parameters in Table 2, as described in Section 4.4.2, without any adjustment to the single-value (i.e. $\sigma_\zeta = 0$ and so $\zeta = 0$ for all individuals). Obtain a set of parameter estimates $\widehat{\Lambda}_0$.
- Step 1. Feed $\widehat{\Lambda}_0$ into the model, and calibrate σ_ζ to match the fraction of singles outside of the main estimation. Obtain $\widehat{\sigma}_{\zeta,1}(\widehat{\Lambda}_0)$.
- Step 2. Draw for each agent a shock ζ from the distribution Z with scale $\widehat{\sigma}_{\zeta,1}$ and add it to their single value, which becomes $w + p^U + \zeta$. Repeat the main estimation, and obtain parameter estimates $\widehat{\Lambda}_1$. Go back to Step 1, but now feed $\widehat{\Lambda}_1$ into the model and calibrate a new σ_ζ to match the fraction of singles outside of the main estimation. Obtain $\widehat{\sigma}_{\zeta,2}(\widehat{\Lambda}_1)$.

Repeat steps 1 and 2 n times until convergence, i.e., $\widehat{\sigma}_{\zeta,n}(\widehat{\Lambda}_{n-1}) = \widehat{\sigma}_{\zeta,n-1}(\widehat{\Lambda}_{n-2})$. Then we set $\sigma_\zeta = \widehat{\sigma}_{\zeta,n}$, draw for each individual a single shock from distribution Z with scale parameter σ_ζ , add it to their economic value from being single, and proceed once more with the main (internal) estimation to check that our estimates $\widehat{\Lambda}$ satisfy $\widehat{\Lambda} \equiv \widehat{\Lambda}_{n-1} = \widehat{\Lambda}_{n-2}$.

Calibration of the Constant in Labor Market Production. We assume that the constant in the production function is not shared between workers and firms, but accrues to the worker in the form of a minimum hourly wage (i.e., the wage of someone with the lowest human capital \tilde{s} who will be matched to the least productive occupation, $y = 0$). Technically, K is the constant of integration in the wage function. We set K to the 5th percentile of wages among workers with weekly labor hours below 25.

Table A.1: Calibrated Parameters

Parameter	Value
Hourly Minimum Wage K	6.32
Variance of the Shock to the Single Value σ_ζ	8.16

D.2 Internal Estimation

D.2.1 Estimation of Complementarity Parameter ρ

We use Lemma 1 to justify our estimation strategy that uses the moment *correlation of spouses' home production hours* to pin down complementarity parameter ρ of home production function p .

Lemma 1 (Affiliation of Spouses' Hours and Home Production Complementarity.). *For a given couple type $\mathbf{s} = (s_m, s_f)$, if their hours $\mathbf{h}^M = (h_f, h_m)$ are affiliated, then p is supermodular.*

Proof. Note that, based on the conditional choice probability for hours worked,

$$\pi_{\mathbf{s}}^M(\mathbf{h}^M) = \frac{\exp(\bar{u}_{\mathbf{s}}^M(\mathbf{h}^M)/\sigma_\delta)}{\sum_{\tilde{\mathbf{h}}^M \in \{\mathcal{H} \cup \emptyset\}^2} \exp(\bar{u}_{\mathbf{s}}^M(\tilde{\mathbf{h}}^M)/\sigma_\delta)},$$

the hours \mathbf{h}^M of spouses of type \mathbf{s} are *affiliated* if $\pi_{\mathbf{s}}^M(h_f, h_m)$ is log-supermodular in (h_f, h_m) . That is, for all $h_f'' > h_f'$ and $h_m'' > h_m'$,

$$\begin{aligned} \frac{\pi_{\mathbf{s}}^M(h_f'', h_m'')}{\pi_{\mathbf{s}}^M(h_f', h_m'')} &\geq \frac{\pi_{\mathbf{s}}^M(h_f'', h_m')}{\pi_{\mathbf{s}}^M(h_f', h_m')} \\ \Leftrightarrow \frac{\exp(\bar{u}_{\mathbf{s}}^M(h_f'', h_m'')/\sigma_\delta)}{\exp(\bar{u}_{\mathbf{s}}^M(h_f', h_m'')/\sigma_\delta)} &\geq \frac{\exp(\bar{u}_{\mathbf{s}}^M(h_f'', h_m')/\sigma_\delta)}{\exp(\bar{u}_{\mathbf{s}}^M(h_f', h_m')/\sigma_\delta)} \end{aligned}$$

which is the case if and only if

$$\bar{u}_{\mathbf{s}}^M(\mathbf{h}^M) = \frac{w(\psi s_f h_f) + w(s_m h_m) + 2p(1 - h_m, 1 - h_f)}{\sigma_\delta}$$

is supermodular in (h_f, h_m) . This in turn is the case if and only if home production function p is supermodular in spouses' time inputs. \square

Remarks on Estimation. First, the affiliation of labor hours (h_f, h_m) implies the affiliation of *home production hours* $(1 - h_f, 1 - h_m)$. Second, the strength of affiliation in spouses' hours is mirrored by the strength of supermodularity of p (i.e., the strength of home production complementarities). Third, given our parameterization of the production function as CES (Section 4.1), the strength of supermodularity of p is reflected in the level of ρ . Fourth, affiliation of hours implies a *positive correlation* of hours. Thus, using the correlation of home production hours in estimation to pin down ρ is a natural choice.

D.2.2 Estimation Results

Table A.2: Targeted Moments

	Model	Data
M1. LFP Female to Male Ratio	0.7489	0.7864
M2. Full Time Work Female to Male Ratio	0.4019	0.3834
M3. LFP Married to Single Ratio, Women	0.8750	0.8556
M4. LFP Married to Single Ratio, Men	1.0004	1.2534
M5. Correlation of Spouses' Home Hours	0.3055	0.3120
M6. Mean Hourly Wage	17.6474	17.6354
M7. Variance of Hourly Wage	52.6180	53.9061
M8. Overall (90-10) Wage Inequality	3.0885	2.9686
M9. Upper Tail (90-50) Wage Inequality	1.7064	1.7271
M10. Correlation between Spouses' Types	0.4436	0.4468
M11. Gender Wage Gap by Effective Type 3	0.1381	0.1155
M12. Gender Wage Gap by Effective Type 4	0.1380	0.1464
M13. Female LFP of Couple Types 3 and 4	0.7461	0.7308
M14. Female LFP of Couple Types 5 and 6	0.8524	0.8071
M15. Female LFP of Single Women Type 3 and 4	0.8146	0.7320
M16. Female LFP of Single Women Type 5 and 6	0.8653	0.8429

Notes: *LFP* stands for Labor Force Participation. Moments are computed as discussed in Online Appendix OE.1. *Types* refer to human capital types, where types 1-6 correspond to columns 1-6 of Table O.13 in the Online Appendix.

E Quantitative Analysis: Inequality Over Time

Table A.3: Estimated Parameters: 1990-1996 versus 2010-2016

		1990-1996		2010-2016	
		Estimate	s.e.	Estimate	s.e.
Female Relative Productivity in Home Production	θ	0.88	0.05	0.82	0.06
Complementarity Parameter in Home Production	ρ	0.01	0.03	-0.54	0.20
Home Production TFP	A_p	37.61	4.49	38.33	3.46
Elasticity of Output w.r.t. s	γ_1	0.40	0.10	0.63	0.05
Elasticity of Output w.r.t. y	γ_2	0.15	0.20	0.18	0.05
Labor Market Production TFP	A_z	39.98	9.77	42.00	2.29
Female Productivity Wedge	ψ	0.78	0.03	0.85	0.02
Marriage Market Preference Shock (scale)	σ_β	0.02	0.08	0.11	0.01

Notes: s.e. denotes standard errors. See Section 4.4.2 for a description of how these standard errors are computed.

Table A.4: Data and Model Moments: 1990-1996 versus 2010-2016

	1990-1996		2010-2016		Data Diff.
	Model	Data	Model	Data	p-value
M1. LFP Female to Male Ratio	0.6353	0.5875	0.7489	0.7864	0.0000
M2. Full Time Work Female to Male Ratio	0.2406	0.3336	0.4019	0.3834	0.0190
M3. LFP Married to Single Ratio, Women	0.8006	0.6343	0.8750	0.8556	0.0000
M4. LFP Married to Single Ratio, Men	1.0077	1.0844	1.0004	1.2534	0.0019
M5. Correlation of Spouses' Home Hours	0.1390	0.1518	0.3055	0.3120	0.0000
M6. Mean Hourly Wage	16.9549	17.0106	17.6474	17.6354	0.0026
M7. Variance of Hourly Wage	35.0394	37.2476	52.6180	53.9061	0.0000
M8. Overall (90-10) Wage Inequality	2.5261	2.3486	3.0885	2.9686	0.0000
M9. Upper Tail (90-50) Wage Inequality	1.5571	1.5841	1.7064	1.7271	0.0000
M10. Correlation between Spouses' Types	0.4037	0.4052	0.4436	0.4468	0.2069
M11. Gender Wage Gap by Effective Type 3	0.1714	0.1657	0.1381	0.1155	0.4080
M12. Gender Wage Gap by Effective Type 4	0.1684	0.1839	0.1380	0.1464	0.2324

Notes: *LFP* stands for Labor Force Participation. Moments are computed as discussed in Online Appendix OE.1. The last column of the table reports the p-value of the hypothesis test that the difference between the data moments across time periods is zero. We use a standard t-test for the difference in means (M6) and a Levene test for the difference in variances (M7). We use a Fisher transformation to construct the test statistic for the differences in correlations (M5 and M10). We use a two-sample Wald test for differences in ratios across periods (M1, M2, M3, M4, M8, M9, M11, M12). To construct the statistic for the Wald tests for the difference in ratios, we use bootstrap techniques for the variance estimation. *Types* refer to human capital types, where types 1-6 correspond to columns 1-6 of Table O.13 in the Online Appendix.

Figure A.1: Inequality Changes Over Time: Detailed Decomposition

