

MARRIAGE MARKET AND LABOR MARKET SORTING

*PAULA CALVO[†]

ILSE LINDENLAUB[‡]

ANA REYNOSO[§]

July 2021

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 on 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 labor sorting. This puts significant downward pressure on the gender wage gap and within-household income inequality, but it fuels between-household inequality. Our estimated model sheds new light on the sources of inequality in today's Germany, and—by identifying important shifts in home production technology 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, the University of Otago, and the University of Chile 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.

[†]Yale University, **Address:** Department of Economics, Yale University, 28 Hillhouse Avenue, New Haven, CT 06520, US. **Email:** paula.calvo@yale.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 is a defining feature of both the labor market and the marriage market and has important implications for inequality. On the marriage market, the matching of partners with similar education impacts both within- and between-household income inequality. Moreover, positive sorting in the labor market between workers and firms or jobs reinforces wage inequality across skills. But even though inequality in economic outcomes results from agents interacting in *both* the marriage and the labor market, how the interplay of the two markets shapes inequality has not been studied before.

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 labor markets further mitigated gender disparities in Germany over the last decades.

Three sets of facts from the German Socioeconomic Panel (henceforth, GSOEP) show a salient relationship between the marriage and the labor market and motivate our analysis. First, as is well documented in the literature, there is positive assortative matching both on spouses' education in the marriage market and between workers' education and jobs' skill requirements in the labor market. Importantly, there is a gender gap in labor market sorting 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 more similar education tend to split their time similarly between market and house work, and conditioning on hours worked, the gender gap in labor market sorting is significantly smaller.

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

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

We focus on a tractable transferable utility (TU) representation of our model and derive two benchmark equilibria that depend on the model’s primitives. Both feature positive assortative matching between workers and jobs in the labor market driven by productive complementarities. The two equilibria, however, differ in household and marriage outcomes depending on properties of the home production function. On the one hand, if home production exhibits complementarity in partners’ time inputs, a *monotone equilibrium* arises, characterized by positive sorting in the marriage market and labor hours that are increasing in both own and partner’s skills. This equilibrium reflects a ‘progressive’ economy with a high frequency of two-earner households and 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 put downward pressure on 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 needs to 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 baseline 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 heterogeneity within each couple type, and a random component to workers’ skills to account for mismatch in the labor market. Second, we parameterize our model allowing

¹We base this assumption on our evidence of a positive impact of hours worked in the labor market on *hourly* wages in the GSOEP, in Figure O3 (Online Appendix) and Table 6, column (3); and also on previous evidence that more hours worked lead to higher productivity—for instance due to reduced coordination costs among co-workers (e.g., Goldin, 2014).

for gender differences in both home and labor market productivity (the latter can also be interpreted as discrimination) that will be disciplined by the data. We show that this model is identified.

We first estimate our model on data from modern Germany—our benchmark estimation, which focuses on West Germany from 2010 to 2016—and find that spouses’ home production hours are complementary. Our model matches key targeted features of the marriage market equilibrium (such as the degree of marital sorting and the high correlation of home hours within couples) and the labor market equilibrium (such as moments of the wage distributions). To further validate the model, we show that it also reproduces 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.

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

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

Finally, we find that changes in both marriage market sorting and labor market sorting—which increased by 10% and 8%, respectively—significantly affected these inequality shifts. If sorting patterns had stayed constant at their 1990 levels, gender inequalities would 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 narrowed gender disparities, since it was predominantly driven by women’s improved labor sorting; this helped them catch up with men’s pay.

2 The 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 (Adda, Dustmann, and Stevens, 2017; Dias, Joyce, and Parodi, 2018; Angelov, Johansson, and Lindahl, 2016; Kleven, Landais, and Søgaard, 2019). Because the wage rate is kept fixed in these papers, any gender gap in pay can only be attributed to *earnings* and not to *hourly wages* (which is what we focus on). In assuming that hours worked affect workers' productivity in the market, we follow more closely the literature that documents significant labor market returns to hours (Aaronson and French, 2004; Gicheva, 2013; Goldin, 2014; Cortés and Pan, 2019; Bick, Blandin, and Rogerson, 2020).² Other work links gender pay gaps to gender differences in preferences for work flexibility (Bertrand, Goldin, and Katz, 2010; Mas and Pallais, 2017; Cubas, Juhn, and Silos, 2019) and to sorting into occupations that require different time inputs (Erosa, Fuster, Kambourov, and Rogerson, 2017). Finally, there is work highlighting the importance of information frictions for gender pay gaps (without considering the marriage market): If employers believe that women have less market attachment 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 has been to study marriage market sorting using structural models. Several papers have investigated how premarital investments in education interact with marriage patterns in a static framework (Fernández, Guner, and Knowles, 2005; Chiappori, Iyigun, and Weiss, 2009) or over the life cycle (Chiappori, Costa-Dias, and Meghir, 2018) and how post-marital investments in a partner's career interact with family formation and dissolution (Reynoso, 2019). Further, 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 (Reynoso, 2019), or different tax systems (Gayle and Shephard, 2019) affect household behavior and marriage sorting.³ Finally, in models with *exogenous* marriage sorting, Fernández and Rogerson

²Bick et al. (2020) find that hourly wages of U.S. men are non-monotone, increasing until 50 hours per week, and then decreasing. Note that in our sample, hardly anyone (<0.3%) works more than 50 hours per week, 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 influential paper, Voena (2015) also focuses on the adoption of unilateral divorce and its effects on household

(2001) analyze the effect of increased marriage sorting on wage inequality, while [Lise and Seitz \(2011\)](#) focus on its effect on between- and within-household consumption inequality.

Like in these papers, marriage market sorting is an important margin in our model. While education is exogenous, 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 impacting labor market sorting, which is different 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, documenting positive assortative matching between workers and firms ([Card, Heining, and Kline, 2013](#); [Hagedorn, Law, and Manovskii, 2017](#); [Bagger and Lentz, 2018](#); [Bonhomme, Lamadon, and Manresa, 2019](#)); or workers and jobs ([Lindenlaub, 2017](#); [Lindenlaub and Postel-Vinay, 2020](#); [Lise and Postel-Vinay, Forthcoming](#)) without taking the marriage market into account. In turn, [Pilossoph and Wee \(2019b\)](#) consider spousal joint search on the labor market to explain the marital premium, but take marriage market sorting as given. Our contribution is to explore how the forces that determine who marries whom shape labor market sorting and pay.

INTERPLAY BETWEEN MARRIAGE AND LABOR MARKET. Our work is most related to a nascent literature on the interplay between marriage and labor markets. This research has focused on the effects of spouses’ joint labor search ([Pilossoph and Wee, 2019a](#); [Flabbi, Flinn, and Salazar-Saenz, 2020](#)) and of 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) and our finding on the key role of home production complementarities/substitutabilities for shaping equilibrium—both in theory and in the data.

3 Descriptive Evidence

3.1 Data

We use two data sources. The German Socioeconomic Panel (GSOEP) is a household panel of around 25,000 individuals (including the surveyed households’ head and the spouse), surveyed yearly. It contains detailed information on labor market outcomes and time use. We focus on West Germany, 2010-2016. In turn, the Employment Survey of 2012 (BIBB) contains detailed occupational characteristics. Details on the GSOEP are in Online Appendix [OC](#) and on the BIBB in Appendix [C.3](#).

behavior, especially asset accumulation. In her paper, the marriage market is exogenous.

⁴Exception 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).

3.2 Empirical Evidence

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

MARRIAGE MARKET SORTING. We find evidence of positive assortative matching (PAM) in education in the German marriage market, in line with previous evidence (Eika et al., 2019 for the US and Germany, and Greenwood et al., 2016 and Greenwood et al., 2017 for the US). Table 1 reports the matching frequencies by education for the period 2010-2016, suggesting that almost 60% of individuals marry someone of the same education level. The correlation between the education level of spouses, which is our main measure of marriage market sorting, is equal to 0.47.⁵ Furthermore, marriage market sorting increased over time. For the period 1990-1996, the correlation between partners’ education was 0.44.

Table 1: Marriage Matching Frequencies by Education

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

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


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

Figure 1 (left) plots the fitted matching function (job attribute as a function of worker characteristic) by gender, conditional on employment (solid lines). Both men and women are positively sorted in the labor market, indicated by a positive slope of the matching function. However, men are ‘better’ matched: For a given education level, men are on average matched to more demanding jobs. This pattern is also reflected in the correlation of worker and job attributes by gender, which is 0.64 for men and 0.62 for women.⁶

⁵As discussed in Chiappori, Dias, and Meghir (2020), *correlation* is one of the measures of sorting that has two desirable properties such a measure should have: a ‘monotonicity’ condition and whether it captures the case of ‘perfectly assortative matching’. Eika et al. (2019) propose an alternative measure of marriage sorting based on the frequency of couples with similar education relative to random matching. This measure equals 1.73 in our sample, 2010-2016: Individuals are 73% more likely to marry someone with the same education, relative to random matching. And during 1990-1996, this sorting measure is 1.59, also suggesting increased sorting over time.

⁶Differences in labor market sorting across gender are even larger (with a correlation 0.58 for men vs. 0.48 for women)

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



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

THE ROLE OF HOURS. We now provide evidence on a salient link between the two markets: hours worked on the labor market vs. hours spent in home production. First, we document that the time allocation choice is related to the partnership status as well as marriage market sorting. Second, we document that at the same time, the time allocation choice is linked to labor market sorting.


As is well documented (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. Among singles (Online Appendix OA.4, Figure O2, left panel), gender differences in time allocation choices are small. But for couples, they are pronounced (right panel). Indeed, in couples, women spend about 12.5 hours less per week working on the labor market but about 20 hours more in home production compared with their male partners. Neither for couples nor for singles are there significant gender differences in leisure, which justifies our decision to abstract from it in our model below.

We also document the relationship between hours and marriage market sorting. Figure 2 shows the correlation between home production hours within couples (left) and the correlation of labor hours

if we do not condition on participation and we treat unmatched individuals as matched to a job with attribute zero.

within couples (right), both against our summary measure of marriage market sorting (difference in partners’ years of education). Interestingly, both home production hours and labor market hours are more complementary between those partners who are well sorted in the marriage market, as indicated by the inverse U-shape of the hours’ correlation functions. Note that the pattern for home production is even more pronounced than for market hours, with a stronger positive correlation of home hours among partners with the same education compared with partners with differences in education.⁷

Figure 2: Time Allocation and Marriage Sorting



We further explore what drives the complementarity in the home production time of spouses by looking at different components of home production, especially since the literature on family economics emphasizes household specialization. We find that complementarities are strongest in childcare, and are least pronounced when it comes to housework (see Figure O1 in Online Appendix OA.2).⁸

One concern is that the relationships between marital sorting and complementarities in hours in Figure 2 are based on marriage market sorting bins that pool individuals from different education groups. Not controlling for education allows for the possibility that hours only depend on own education but do not vary with partner’s education if, e.g., low (high) educated workers always supply low (high) hours independent of the partner’s type. Also, the complementarities in partners’ hours from Figure 2 might be driven by confounding factors or unobserved heterogeneity that drive both spouses’ hours choices. We discuss and address several of these concerns in Online Appendix OA.3. There, we control for education and other covariates that might be correlated with the choice of hours (Table O2). Moreover, we show that partners’ complementarities in labor market hours are stronger when we exploit exogenous variation in childcare availability across states and time to instrument for female labor market hours (Table O3). Finally, we also show in these regressions that the correlation between partners’ hours is larger for those who are better sorted in the marriage market, in line with the descriptive evidence of

⁷For consistency between the right and left panels of Figure 2, we condition on both partners participating in the labor market. The pattern in the left panel also holds if we do not condition on labor market participation.

⁸In recent work, Moschini (2021) also finds strong complementarities between parental time inputs when it comes to investment in children’s skills.

Figure 2. We provide the details of our identification strategy in Online Appendix OA.3.

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

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

SUMMARY. We highlight three sets of facts. First, there is evidence of PAM in both the labor and the marriage market. In the labor market, however, men are ‘better’ matched than women. Second, there is a strong relation between labor market and marriage market sorting, with labor market sorting being maximized when marriage market sorting is. Third, the split between hours worked in the labor market vs. hours spent in home production is a potentially important link between the two markets. We not only show that time allocation choices depend on marriage market sorting but also that they are themselves associated with labor market sorting. Motivated by these facts, we now build a model with endogenous labor *and* marriage markets, linked through households’ labor supply choices. We come back to these facts when validating our model below.

4 The Model

We first lay out the environment and agents’ decisions, and then define equilibrium.

4.1 Environment

There are two types of agents, individuals and firms. There is a measure one of firms. Firms have productivity attribute $y \in \mathcal{Y} = [y, \bar{y}]$, distributed according to a continuously differentiable 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: Denote women’s skills by $x_f \in \mathcal{X}_f = [0, \bar{x}_f]$, where x_f is distributed with a continuously differentiable cdf N_f , with density $n_f > 0$. Analogously, men have skills $x_m \in \mathcal{X}_m = [0, \bar{x}_m]$, 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, what matters for output is not only skills but also hours worked, which will be chosen optimally in each couple. Each individual is endowed with one unit of time that can be allocated to paid work in the labor market, denoted by $h_i, i \in \{f, m\}$, or non-paid work at home toward the production of a public good, $\ell_i = 1 - h_i$ (based on Figure O2, Online Appendix, which shows no large differences in leisure across gender, we abstract from it). Note that $h_i = 0$ captures non-participation. By increasing hours worked in the labor market, each individual ‘invests’ in his/her *effective skill* $\tilde{x} := e(x, h)$, $\tilde{x} \in \tilde{\mathcal{X}}$, with endogenous cdf $\tilde{N}(t) := \mathbb{P}[\tilde{x} \leq t] = \frac{1}{2}\mathbb{P}[\tilde{x}_f \leq t] + \frac{1}{2}\mathbb{P}[\tilde{x}_m \leq t]$. 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. The effective skill or index \tilde{x} is the output-relevant worker characteristic on the labor market.⁹ This 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 generated by an individual of type \tilde{x} matched to a firm of type y . We assume that production function z is twice continuously differentiable and strictly increasing in each argument. Individuals and firms split the output they generate into wages and profits, where workers use their wages to finance the private consumption good $c_i, i \in \{f, m\}$.

The public good production function is denoted by p . 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 spending $(\ell_m, \ell_f) = (1 - h_m, 1 - h_f)$ in home production (recall that hours at home equal the hours *not* spent working).¹⁰ We assume that p is twice continuously differentiable with $p_{\ell_m} > 0, p_{\ell_f} > 0, p_{\ell_m \ell_m} < 0$ and $p_{\ell_f \ell_f} < 0$ and satisfies the Inada conditions $\lim_{h_f \rightarrow 0} p_{\ell_f}(1 - h_m, 1 - h_f) = 0$ and $\lim_{h_f \rightarrow 1} p_{\ell_f}(1 - h_m, 1 - h_f) = \infty$, and similarly for p_{ℓ_m} .¹¹

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

Both matching markets—the labor and marriage market—are competitive (full information and no

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

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

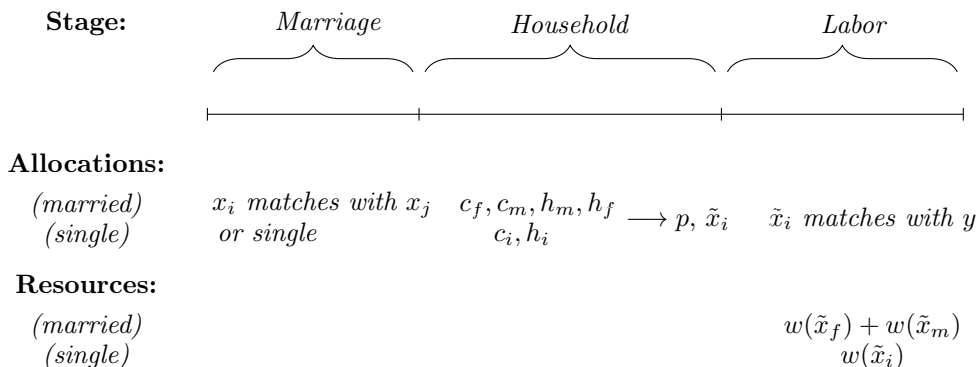
¹¹Throughout, we will denote the partial derivatives of some generic function $f(x, y)$ using subscripts; for instance, the derivative of f with respect to x is f_x unless there is risk of confusion, in which case we use $\partial f / \partial x$. We will denote the derivative of a function of a single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

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

4.2 Decisions

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

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



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

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

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

This problem, along with market clearing, pins down the labor market matching function $\mu : \tilde{\mathcal{X}} \rightarrow \mathcal{Y}$, mapping workers’ effective skills to firm types in a measure-preserving way.¹² Matching function μ depends on complementarities of (\tilde{x}, y) in z . Importantly, it also depends on the hours choice (through

¹²Since we focus on monotone matching below, we restrict attention to pure matching, given by a function μ .

\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 below—then the first-order condition of problem (1), which gives a differential equation for w , pins down the wage function as

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

where w_0 is the constant of integration.¹³

HOUSEHOLD PROBLEM. Consider a couple (x_f, x_m) who anticipates (w, μ) in the labor market. This couple solves the following *cooperative household problem*. One partner (here w.l.o.g. 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:

$$\begin{aligned} \max_{c_m, c_f, h_m, h_f} \quad & u(c_m, p(1 - h_m, 1 - h_f)) \\ \text{s.t.} \quad & c_m + c_f - w(\tilde{x}_m) - w(\tilde{x}_f) \leq 0 \\ & u(c_f, p(1 - h_m, 1 - h_f)) \geq \bar{v}, \\ & 0 \leq h_i \leq 1, i = f, m \end{aligned} \quad (3)$$

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

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

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

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

¹³Note that $w(\tilde{x})$ is the wage of a worker with effective type \tilde{x} *per unit of time* (recall that our time endowment is normalized to one unit and we will give a specific interpretation of a time unit in the data below), not the worker's earnings.

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

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

4.3 Equilibrium

We now formally define equilibrium.

Definition 1 (Equilibrium). *An equilibrium is given by a tuple 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 below.

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

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

In a monotone equilibrium, there are three additional requirements relating to the three different stages of this model. 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 , implying that γ_f can be inverted; and similarly for men with $\gamma_m(x_m)$. As a result, in a monotone equilibrium, the endogenous cdf of effective types can be written as follows:¹⁴

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

This discussion highlights an important point. The equilibrium hours function, h_f , not only depends on her own skill type x_f but also on marriage market outcomes: the skill type of her partner, $\eta(x_f)$, as well as the transfer guaranteed to her in the marriage $v(x_f)$; and similarly regarding the factors the male

¹⁴To see this, observe that the probability that $\tilde{x} \leq t$ is:

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

hours function h_m depends on. As a result, labor supply choices form the link between the marriage 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 challenging from a theoretical point of view, since we must simultaneously equilibrate two intertwined matching markets, which are related through the time allocation choices.

5 Analysis

In this section we show how the primitives of our model—in particular, home production p and labor market production z —shape equilibrium. To gain tractability, we focus on a certain class of models (the quasi-linear class) that yields the *transferable utility* (TU) property. This circumvents a further potential complication of *imperfectly transferable utility* (ITU), in which the hours functions and thus the public good production depend on transfer v .

5.1 The Quasi-Linear Class

The TU representation of our model obtains if the utility function falls into a known class, the *Gorman form*, which guarantees that utility is linear in the private good (possibly after a monotone transformation); see also [Mazzocco \(2007\)](#); [Browning et al. \(2014\)](#).¹⁵ We assume the utility function:

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

which has the important property of being quasilinear in c_i .¹⁶ Then 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, in turn, the couple’s marital surplus is independent of the sharing rule v . As a result, the matching problem can be solved by maximizing the total value of marriage, independent of how it is shared (as in [Shapley and Shubik, 1971](#) and [Becker, 1973](#)).

5.2 Conditions for Monotone Equilibrium

Our objective is to derive conditions under which any equilibrium that satisfies certain basic regularity properties is *monotone* in the sense of Definition 2. To summarize these regularity properties in a compact way, we first define *regular equilibrium*.

¹⁵More generally, the Gorman form of 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 $k(p)$, meaning that the marginal utility w.r.t. c_1 is equalized across partners, so that utility can be transferred between them at a constant rate.

¹⁶We here choose the simplest functional form in the Gorman class to reduce notation that obscures the main mechanism. In Online Appendix [OB.1](#) we provide examples of other utility functions in this class under which our model features TU.

Definition 3 (Regular Equilibrium). *An equilibrium is regular if household problem (3) has a unique solution that is interior and continuous in (x_m, x_f) , and \tilde{N} is atomless.*

In what follows, we will focus on equilibria that are regular. We detail in Appendix A.2 why these regularity assumptions are important. At the cost of much additional technical detail, we could justify the regularity properties of Definition 3 in terms of primitives.¹⁷ The approach we follow, however, allows us to focus on the monotone structure of the model in a more direct way. The monotone equilibrium will be our benchmark. We first state our main result and we then unpack its components.

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 is in Appendix A.2. A crucial condition for the monotone equilibrium is the complementarity of spousal time in home production (supermodular p , $p_{\ell_m \ell_f} > 0$), which gives rise to a ‘progressive’ way of organizing the household with gender balance in hours as opposed to specialization. In this case, if an individual puts many hours into labor market work then the partner will do the same, at the cost of less home production. The positive correlation of partners’ hours within the household is clearly a force toward PAM in the marriage market: Having a partner with similar skills makes it easier to work similar hours in the labor market and, on the flip side, to put similar hours into home production, reaping the full benefits from home production complementarity. In turn, positive sorting in the labor market stems from the complementarity between individuals’ effective skills and jobs’ skill requirements (supermodular z , $z_{\tilde{x}y} > 0$).

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

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

HOURS ARE INCREASING ON OWN AND PARTNER’S TYPE. With quasi-linear utility (6), the household problem (3) takes the form:

$$\begin{aligned} \max_{h_m, h_f} w(\tilde{x}_m) + w(\tilde{x}_f) - \bar{v} + 2p(1 - h_m, 1 - h_f), \\ \text{s.t.} \quad 0 \leq h_i \leq 1, \quad i = f, m \end{aligned} \tag{7}$$

where we substituted both the household’s budget constraint and the wife’s constraint to receive at least utility \bar{v} into the objective function. As a consequence of TU, the overall split between public

¹⁷Here we provide a sketch: The property of interior solutions of the household problem can be justified based on Inada conditions on p . In turn, we can specify conditions on the objective function of the household problem that render a unique solution (which also implies stability of the equilibrium in a tatonnement sense). Uniqueness, in turn, allows for the application of the Implicit Function Theorem, which guarantees differentiability and therefore continuity of the hours functions in (x_m, x_f) . Finally, an atomless \tilde{N} can be guaranteed if, in addition, the effective type function e is assumed to be Morse (i.e., a function with only isolated critical points).

and private consumption can be made independently of how utility is shared, captured by \bar{v} . As a consequence, the hours functions (h_f, h_m) will only depend on types (x_f, x_m) but no longer on \bar{v} .

The FOCs of household problem (7) with respect to h_f and h_m at an interior solution are given by:

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

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

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

To characterize under which conditions hours are increasing in own and partner's type, we apply a monotone comparative statics argument. In Appendix A.2, we show that the hours functions feature the discussed monotonicity in both skills if: (i) home production hours are complementary (p is supermodular) and if (ii) wages are supermodular in skills and hours, $(w)_{x_i h} = w''e_{x_i}e_h + w'e_{x_i h} > 0$ where $h = h_i(x_m, x_f)$ (guaranteed by convex and supermodular z , which renders the wage function convex). We show that hours are strictly increasing in partners' skills at any (x_m, x_f) and thus also along the marriage market equilibrium assignment $(\eta(x_f), x_f)$.¹⁸

POSITIVE SORTING IN THE MARRIAGE MARKET. Given the equilibrium hours functions (h_f, h_m) , we obtain the value of marriage Φ as the value of household problem (7):

$$\begin{aligned} \Phi(x_m, x_f, v(x_f)) = \\ w(e(x_m, h_m(x_m, x_f))) + w(e(x_f, h_f(x_m, x_f))) - v(x_f) + 2p(1 - h_m(x_m, x_f), 1 - h_f(x_m, x_f)). \end{aligned} \quad (10)$$

Complementarities among partners' types in Φ determine marriage market matching patterns. Under TU, $(\Phi)_{x_m x_f} = \Phi_{x_m x_f}$. If $\Phi_{x_m x_f} > 0$, then marriage matching is PAM, and thus $\eta' > 0$.¹⁹

Using once more a monotone comparative statics argument, we show that Φ is supermodular if wages are complementary in types, $(w)_{x_m x_f} > 0$. Zooming in, this complementarity in types derives from wages that are supermodular in type and hours, $(w)_{x_f h} > 0$ and $(w)_{x_m h} > 0$, in combination with labor hours that are increasing in the partner's type—a property that holds in monotone equilibrium.

These sorting conditions are intuitive: There is PAM in the marriage market, so that x_f is matched to $\eta(x_f) = N_m^{-1}(N_f(x_f))$, if the labor hours of spouses are complementary in the sense that an individual's labor hours are increasing in partner's type and if, at the same time, working more hours boosts the marginal wage return to skill. In turn, the property that own labor hours are increasing in the partner's

¹⁸An additional important consequence of the strict monotonicity of the hours functions is that the distribution of effective types, \tilde{N} , can be pinned down and is given by (5). Moreover, \tilde{N} is atomless and \tilde{X} is an interval in this case.

¹⁹Under ITU (Legros and Newman, 2007), $(\Phi)_{x_m x_f} > 0$ is equivalent to $\Phi_{x_m x_f} > \frac{\Phi_{x_f}}{\Phi_v} \Phi_{x_m v}$ (Chade, Eeckhout, and Smith, 2017). To see this, plug the FOC of problem (4), $\Phi_{x_f} + \Phi_v v' = 0$, into $(\Phi)_{x_m x_f} = \Phi_{x_m x_f} + \Phi_{x_m v} v'$. In the quasi-linear class, the condition for PAM simply becomes $\Phi_{x_m x_f} > 0$ since $\Phi_{x_m v} = 0$.

type crucially derives from the home production complementarity, as we have shown above. This highlights the importance of the home production function also at the marriage stage.²⁰

5.3 Properties of Monotone Equilibrium and Stylized Facts

We now connect the properties of monotone equilibrium with our stylized facts from Section 3 in a *qualitative* way, before accurately replicating our facts in our quantitative analysis below.

MARRIAGE MARKET SORTING. The property of positive sorting in the marriage market resembles our empirical finding of positive sorting on partners' education in Table 1.

LABOR MARKET SORTING. In a monotone equilibrium, more skilled individuals work more in the labor market than at home compared with the less skilled, a feature that is reinforced by having more skilled partners. As a result, more skilled individuals have higher effective types, and thereby obtain more productive labor market matches: There is positive sorting in the labor market in (x, y) , which captures the positive correlation between education and jobs' skill requirements in the data (Figure 1, left).

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

$$(\mu)_{x_f} = \mu' \left(e_{x_f} + e_h \left(\frac{\partial h_f}{\partial x_m} \eta' + \frac{\partial h_f}{\partial x_f} \right) \right), \quad (11)$$

which is based on $\mu(\tilde{x}_f)$ with $\tilde{x}_f = e(x_f, h_f(\eta(x_f), x_f))$. Equation (11) illustrates how labor market sorting μ depends on marriage market sorting η . When marriage market sorting is positive, $\eta' > 0$, then higher x are matched to higher y , $(\mu)_x > 0$ (given that the hours of spouses are complementary $\partial h_f / \partial x_m > 0$). The intuition is straightforward. PAM in the marriage market induces individuals with higher x_i to have a better partner $x_j = \eta(x_i)$ and therefore to work more hours, which translates into a higher effective type \tilde{x}_i and thus a better labor market match $y = \mu(\tilde{x}_i)$, compared with when marriage market sorting is not positive. In a stylized way, this property of the monotone equilibrium is related to our empirical fact that labor market sorting is stronger for positively sorted couples (Figure 1, right).

THE ROLE OF HOURS. In a monotone equilibrium, labor hours are complementary within couples: Increasing, say, female skills, not only pushes up her own labor hours but also induces her partner to

²⁰Analogous to the labor market, we can complete the equilibrium construction in the marriage market by obtaining the market-clearing price v . Maximizing (10) with respect to x_f , while taking into account that $h_m(x_m, x_f)$ and $h_f(x_m, x_f)$ are already optimized so that they do not respond to further changes in x_f (by the envelope theorem), yields:

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

which will be evaluated at η . The transfer to the female partner, v , reflects the marginal impact of her type on her wage: When a man chooses a woman, he trades off the marginal benefits of choosing a higher type (which equals the marginal impact on her wage, $w' e_{x_f}$) with the marginal costs (which equals the marginal increase in transfer to her, v').

work more. As a result, partners' hours comove. There are two drivers behind this result. First, for a given male partner type x_m (for *exogenous* marriage matching), an increase in female skills increases her labor hours. But this reduces her home hours, which induces her partner to also work less at home and more in the market due to the home production complementarity, $p_{\ell_m \ell_f} > 0$. As a result, *both* partners increase their labor hours as the female skill improves. Second, this complementarity in hours is reinforced under *endogenous* marriage market sorting: Under PAM, an increase in her skill x_f leads to a better partner $x_m = \eta(x_f)$, who by himself puts in more labor hours and fewer home hours. And since $p_{\ell_m \ell_f} > 0$, the wife adjusts hours in the same direction (fewer home hours and more labor hours), reinforcing the comovement of hours within the couple. Thus, PAM in the marriage market fuels the complementarity of hours within couples—a feature we observe in the data (Figure 2).

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

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

5.4 Non-Monotone Equilibrium

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

Proposition 2 (Non-Monotone Equilibrium). *Assume p is strictly submodular 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 non-monotone.*

The proof is in Appendix A.3. This result highlights the crucial role of home production complementarities/substitutabilities in shaping equilibrium. Making hours at home substitutable, $p_{\ell_m \ell_f} < 0$, gives rise to an equilibrium that relies on ‘specialization’, in which a more skilled partner puts in more labor hours while own labor hours go down in response. At the same time, the partner spends less time in home production, while own home production time increases. This specialization within the household is clearly a force 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. Skilled individuals then prefer to match with less skilled partners.

The only feature that both equilibria have in common is PAM in the labor market, not only in (y, \tilde{x}) (guaranteed by supermodular z) but, importantly, also in (y, x) , which follows from equation (11).²¹

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

6 Quantitative Model

To evaluate the quantitative importance of our main insights in Section 5, we now augment our model so that it can match the data while preserving its mechanism.

6.1 Setup and Decisions

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

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

²¹Further note that the distribution of effective types $\tilde{N}(t) = \frac{1}{2}N_f(\gamma_f^{-1}(t)) + \frac{1}{2}N_m(\gamma_m^{-1}(t))$ will be pinned down as in a monotone equilibrium, since own labor hours are still increasing in own type so that γ_i is still invertible.

productivity y now solves: $\max_{\tilde{s}} z(\tilde{s}, y) - w(\tilde{s})$.

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

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

$$\begin{aligned} & \max_{c_m, c_f, h_m, h_f} u(c_m, p^M(1 - h_m, 1 - h_f)) + \delta^{h_m} \\ \text{s.t.} \quad & c_m + c_f - w(\tilde{s}_m) - w(\tilde{s}_f) = 0 \\ & u(c_f, p^M) + \delta^{h_f} \geq \bar{v}, \end{aligned} \tag{12}$$

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

$$\begin{aligned} & \max_{c_i, h_i} u(c_i, p^U(1 - h_i)) + \delta^{h_i} \\ \text{s.t.} \quad & c_i - w(\tilde{s}_i) = 0, \end{aligned} \tag{13}$$

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

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

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

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

Similar to the baseline model, Φ captures the *economic* surplus from marriage. Different from the baseline model, due to the introduction of labor supply shocks that have not yet realized at the time of marriage, Φ is the *expected* economic surplus from marriage. The expectation is taken over the different hours alternatives of the couple whose choice probabilities are pinned down at the household stage (details are in Appendix B). 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 $\sum_{(s_m, s_f) \in \{\mathcal{S} \cup \emptyset\}^2} \eta_{(s_m, s_f)} = 1$.

6.2 Functional Forms

We parameterize our model as follows. 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, (γ_1, γ_2) are the curvature parameters reflecting the elasticity of output with respect to effective human capital and firm productivity, and K is a constant.

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

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

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

The utility function of individual i is given by $u(c_i, p) = c_i + p$, where $p \in \{p^M, p^U\}$ for spouses and singles, and where we assume that both men and women have the same preferences. We adjust the private consumption of singles by the McClemons equivalence scale (Anyaeibu, 2010).²²

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

We then specify the effective human capital functions as

$$\begin{aligned} \tilde{s}_f &= \psi s_f h_f \\ \tilde{s}_m &= s_m h_m, \end{aligned} \tag{14}$$

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 gender discrimination or productivity differences.

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

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

where we allow for different preference shock distributions for marriage partners and singles and index t indicates the household type.²³ We normalize the location parameter of both labor supply and marriage

²²The McClemons equivalence scale is a standard adjustment of the consumption of singles so that it can be compared to that of married individuals who enjoy economies of scale in consumption.

²³Note that without the different scales for partner and single choices, our parsimonious model (featuring no couple/single-specific parameters) would have difficulty generating enough singles. Allowing for different scales, however, means that our marriage market resembles a nested logit problem with degenerate (single) nest, associated with known identification issues for the scale of the degenerate nest (Hunt, 2000). This is why we fix σ_β^U outside of the main estimation below.

market preference shocks to zero, $\bar{\delta} = \bar{\beta}^M = \bar{\beta}^U = 0$. Moreover, we specify the labor supply shocks as

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

That is, when making hours choices, a decision-making unit draws *a single* labor supply shock, δ^{h^t} , that is extreme-value distributed. In the case of singles, the decision-making unit is just one person and hence, as is standard, this agent draws a shock for each hours alternative. In the case of spouses, however, the decision-making unit is the couple. Therefore, a household draws a single shock for each *joint* time allocation of the spouses (equivalently, the sum of the spouses' shocks is assumed to be extreme-value distributed).²⁴ We make this adjustment to the standard setting, where *each* individual agent draws an extreme-value shock when making a discrete choice, in order to obtain tractable choice probabilities for the joint hours allocation that help with computation and identification of the model.²⁵

6.3 Model Solution

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

A challenge in our fixed-point algorithm is that when partners determine whether a particular hours choice is optimal (which, as discussed, can be understood as an ‘investment’ in effective skills), they must compare the payoff of this investment with all alternative investments. But the competitive wage only determines the price for *equilibrium* investments.²⁶ In order to obtain the off-equilibrium wages without significantly perturbing the equilibrium wages, we use a *tremble strategy*. We postulate that a small fraction of agents are tremblers who make a mistake by choosing off-equilibrium hours. This ensures that off-equilibrium choices will be also priced and individuals can compare *all* investment choices when solving the household problem. While trembling is a widely used concept in game theory, we believe the application to matching markets with investment is new.

²⁴Given the assumed quasi-linear utility function, the household problem depends on the *sum* of spouses' shocks.

²⁵Gayle 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.

²⁶This issue is similar to the one in Cole, Mailath, and Postlewaite (2001), who study bargaining in a matching problem with pre-match investment. Apart from trembling, the hours shocks also help us here to price all investment alternatives.

7 Estimation

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

7.1 The Data

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

7.2 Identification

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

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

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

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

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

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

Proposition 3 (Identification). *Under the assumed functional forms, the model parameters are identified.*

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

7.3 Two-Step Estimation

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

7.3.1 First Step: Estimation Outside the Model

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

ESTIMATION OF WORKER TYPES. Based on our theory, we specify an empirical model for hourly wages as a function of effective types (which in turn are a function of education, ability, and hours worked). We then make use of the longitudinal structure of the GSOEP to estimate workers' unobserved heterogeneity ν from individual fixed effects in a panel wage regression. For computational tractability, we divide individuals in each education bin $ed \in \{hs, voc, c\}$ (standing for high school or less, vocational training and college, respectively) into two groups depending on their ν (above or below the median). Hence, individuals belong to one of six human capital bins (three education types times two ability types). We then order individuals by their human capital $s_i = \alpha^{ed} x_i^{ed} + \nu_i$ (where x_{it}^{ed} are indicator variables for the education group of an individual, meant to capture x in our model, and α^{ed} are the estimated returns to education in the panel wage regression). This gives us a global ranking of worker types. We use the empirical cdf of s_i as our estimate for workers' human capital distribution N_s .

We address three challenges in implementing the panel wage regression: First, we use an instrumental variable (IV) approach to account for the endogeneity of worked hours. Second, we apply a Heckman selection correction (Heckman, 1979) to account for non-random selection in labor force participation. Third, we impute the fixed effects for those individuals for whom we cannot estimate their human capital types, since they appear less than two periods in our panel.

We provide details on this estimation of worker types, including the IV, selection, and imputation in Appendix C.2. The estimation results are in Table 6 and the estimated skill distribution in Table 7. Note that our panel wage regression—under the IV approach—delivers a causal effect of worked hours on hourly wages, where we find that increasing weekly hours worked from 30 to 40 increases hourly wages by around 4%.²⁷

ESTIMATION OF JOB TYPES. The empirical counterpart of our model's firms are occupations (we do not observe firms in the GSOEP). As in Section 3, we measure occupations' productivity y from data on their task complexity. Our main dataset is the BIBB (comparable to the O*NET in the US)

²⁷This estimated effect in our sample of men and women in Germany is smaller but comparable to effects estimated on US data: Aaronson and French (2004) also use a panel regression with fixed effects with an IV for hours and find that increasing hours from 20 to 40 per week increases the hourly wage by 25%; Bick et al. (2020) (who focus on men) find that increasing hours from 30 to 40 per week increases the hourly wage by 11%.

that contains extensive information on tasks performed in each occupation. We focus on 16 tasks measured on a comparable scale. We measure the occupations’ types in two steps. First, we use a Lasso wage regression to select the important/pay-off relevant tasks. In a second step, we run a principal component analysis (PCA) to reduce the task dimensions further to a single one, in which we use the (normalized) first principal component as our one-dimensional occupation characteristic y . Importantly, we use the wage regression only to select the relevant tasks, but we do not use the estimated coefficients. In Appendix C.3, we provide the details of this approach and alternative approaches we pursued for robustness and which yielded similar results.

ESTIMATION OF THE CONSTANT IN PRODUCTION FUNCTION. We assume that the constant in the production function is not shared between workers and firms, but accrues to the worker in the form of a minimum hourly wage (i.e., the wage of someone with the lowest human capital who will be matched to the lowest productive occupation, $y = 0$). In this way, we obtain $K = 6.32$.

7.3.2 Second Step: Internal Estimation

There are nine remaining parameters of the model, $\Lambda \equiv (\theta, \rho, A_p, \gamma_1, \gamma_2, A_z, \psi, \sigma_\delta, \sigma_\beta^M)$. They are disciplined by 17 moments that we chose based on our identification arguments (Section 7.2). To estimate these parameters, we apply the method of simulated moments (McFadden, 1989; Pakes and Pollard, 1989). For any vector of parameters, Λ , the model produces the 17 moments, $mom_{sim}(\Lambda)$, that will also be computed in the data, mom_{data} . We then use a global search algorithm to find the parameter values that minimize the distance between simulated and observed moments. Formally, the vector $\hat{\Lambda}$ solves

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

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

7.4 Results and Fit

We report the parameters we fixed outside of the structural estimation, $(K, \bar{\delta}, \bar{\beta}^M, \bar{\beta}^U, \sigma^U)$, in Table 10 (Appendix C.4.1). The estimated parameters are in Table 2. The result we want to highlight is that our estimates indicate that spouses’ hours at home (and therefore in the labor market) are *complements* with $\rho = -0.54$, giving an elasticity of substitution of spousal home production inputs of 0.65 and pushing the model toward the *monotone equilibrium* of the model. The main data moment calling for a negative ρ is the strong positive correlation of spouses’ home hours.²⁸ Further, the estimated home production function indicates that women are significantly more productive at home than men ($\theta = 0.78$). The large differences in labor force participation and full-time work across genders call for this relatively high female productivity at home. In terms of labor market production, our estimates indicate that it is

²⁸Further, if $\rho > 0$ (i.e., if hours were strategic substitutes), then marriage market sorting would be random; see Figure 7d, where marriage sorting drops significantly as ρ becomes positive.

concave in both the workers’ effective human capital as well as the jobs’ productivity ($\gamma_1 < 1, \gamma_2 < 1$).²⁹ Labor market TFP A_z is estimated to be higher than home production TFP, A_p . The empirical gender wage gap *conditional* on hours and human capital calls for a female productivity/discrimination wedge, which we estimate as $\psi = 0.84$. This implies that for any given type and choice of hours, women’s effective human capital is 16% lower than that of men.

Finally, regarding the marriage preference and labor supply shocks, our estimated scale parameters ensure that we match the fraction of singles, the extent of mismatch in the marriage market, and the heterogeneity in hours choices by households of the same human capital type. We also report the standard errors of the estimates.³⁰ The last column presents our sensitivity analysis (Andrews, Gentzkow, and Shapiro, 2017), in which we report the most important moments that explain 50% of the impact on each parameter in estimation.³¹ Our sensitivity analysis is in line with our identification arguments. For example, the correlation of spouses’ hours, M_5 , is an important moment that disciplines the home production complementarities, ρ ; or the female productivity wedge, ψ , is most related to the within-type gender wage gaps, M_{12} and M_{13} .

Table 2: Estimated Parameters

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

Notes: s.e. denotes standard errors. *Top Sensitivity Moments* reports the most important moments explaining 50% of the total impact on each parameter in estimation, based on our sensitivity measure (see footnote 31). M_1, \dots, M_{17} denote the 17 targeted moments (see Table 11, Appendix C.4).


Figure 4 summarizes the fit between model and data moments, in which we plot all 17 moments (red dots indicate the level of these moments in the model) and their blue confidence interval of the corresponding data moment (computed from a bootstrap sample). We rescaled some moments ($M_6 - M_9$) to be able to plot them all in the same graph. Table 11 in Appendix C.4 reports the fit in detail and indicates the moments corresponding to numbers M1-M17. 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.

²⁹Convexity of z is only sufficient but not necessary for monotone equilibrium.

³⁰The 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 10×17 matrix of the partial derivative of moment conditions with respect to each parameter evaluated at $\Lambda = \hat{\Lambda}$, C is the covariance matrix of the data moments, and \mathcal{V} is the weighting matrix used in estimation.

³¹We 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 30 for notation.

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



7.5 Model Validation

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

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

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


RELATIONSHIP BETWEEN LABOR MARKET SORTING AND MARRIAGE MARKET SORTING. We documented, in Section 3.2, a strong link between labor market and marriage market sorting in the data, 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.

Note that consistent with our quantitative model, here we proxy marriage market sorting by spouses’ differences in *human capital s*-types (as opposed to differences in education that we used in Section 3.2), also in the data. Similarly, labor market sorting is measured by the correlation of (s, y) (instead of (x, y)).

HOURS AS THE LINK BETWEEN MARRIAGE MARKET SORTING AND LABOR MARKET SORTING. A major feature of our model is that marriage and labor markets are linked in equilibrium; namely, through the household’s time allocation choice. Here we show that the model replicates salient features of the data, according to which hours are associated with both marriage and labor 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 $s_f \approx s_m$, around the vertical line at ‘zero’). This is a natural prediction of our model: Spouses of similar human capital can better act on the hours complementarity in home production and better align their hours relative to couples with large human capital differences who tend to specialize.

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


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



8 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 focusses on two contexts: Germany in a recent cross-section (Section 8.1) and over time (Section 8.2).

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



8.1 Inequality Through the Lens of our Model

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

8.1.1 Inequality in the Data and the Model

To assess the extent of inequality in data and model, we focus on four measures: the gender wage gap and household income variance, including its decomposition into between- and within-household components. These statistics are reported in Table 3.³² While our model underestimates the level of the income variance (83 in the model versus 98 in the data), we capture the split of within- and between-household inequality quite well (54-46 split in the model vs. 50-50 in the data). Moreover, regarding inequality within households, the model captures the gap between men and women accurately: It predicts that the share of female wages in overall household wage income is 31% (in the data, it is 33%). Last, our model produces a sizable unconditional gender wage gap (23%), slightly overestimating the observed gap (20%).

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

³²The gender wage gap is computed as the difference in mean wages of men and women over men’s mean wage. The within-component of household income is measured by the variance of wages within a couple, averaged across all couples. The between-component is measured as the variance of the average income of each couple. Our measure of the gender wage gap includes all individuals in the sample, singles and in couples, conditional on employment. In turn, both the female’s share of household income and the total income variance and its decomposition are computed based on the sample of couples. All couples are included, independent of employment status.

Table 3: Gender and Household Inequality

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

8.1.2 Comparative Statics with Respect to Key Drivers of Inequality

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

We can think of several policies and technological changes that impact these parameters. Anti-discrimination policies (such as gender quotas or equal pay policies) can affect ψ . Childcare availability and parental leave policies (such as “daddy months”) might affect θ . Since the child-related tasks women were expected to perform at home are performed by someone else, gender differences in productivity of home hours are likely to decline. Further, changes in home production technologies that facilitate the house chores women traditionally specialized in (Greenwood et al., 2016; Greenwood, 2019) can affect θ and ρ . Finally, an increase in the returns to parents’ joint investment in children (Ramey and Ramey, 2009; Lundberg, Pollak, and Stearns, 2016) could also impact ρ .

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

Figure 7, first row, plots the effect of ρ on different inequality measures: the gender wage gap (panel a), within- and between-household income inequality (panel b), and overall household income inequality

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


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

THE EFFECT OF θ . Next, we are interested in the effect of women's relative productivity at home on our inequality measures. Eliminating the gap in home productivity (reducing θ from the estimate $\theta = 0.78$ to $\theta = 0.5$) would cut the gender wage gap by almost half (Figure 12a in Appendix D.1). The mechanism is as follows: Making female and male home productivity more equal increases the incentive for positive marriage sorting (panel d) and pushes away from household specialization with a smaller gender gap in labor hours (panel e). This positively affects women's wages directly and also indirectly, through a smaller labor market sorting gap (panel f). Interestingly, overall household income inequality decreases as men and women become similarly productive at home, Figure 12c. Here, this is driven by a decrease in within-household inequality (mimicking the evolution of the gender wage gap), which dominates the increase in between-household inequality driven by a rise in marriage sorting (Figure 12b).

In Appendix D.1, we also analyze in detail the comparative statics of the female labor market wedge ψ . Eliminating that wedge (increasing ψ from our estimate $\psi = 0.84$ to $\psi = 1$) has qualitatively very similar effects to reducing θ .

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

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



8.2 Inequality Over Time


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

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

To assess over-time changes in inequality with our model, we compare our estimation for 2010-2016 with the re-estimated model in an earlier period, 1990-1996.³³ For re-estimation on the 1990-1996 sample, we re-assess the skill and job distributions for the earlier period and re-estimate all parameters except those pertaining to the labor supply preference shock, which we set to the level of our current period benchmark (Section 7.4). This is to tie our hands and force the model mechanism to explain the

³³In the GSOEP, 1990 is the first year that features the time-use variables used in our analysis of the later period.

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



data, as opposed to giving changes in shock distributions a too prominent role.³⁴ The model fit along targeted moments is shown in Table 13 in Appendix D.2, which also indicates that both labor and marriage market underwent statistically significant changes over time (column 5). Regarding the untargeted inequality moments of the data (Figure 8, left panel), the model replicates the over-time changes quite well, in which the turquoise bars indicate changes in the data and the purple bars changes in the model.

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


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

In line with our comparative statics exercises, the documented changes in home production technology reduced gender disparities (gender wage gap and within-household inequality) as well as overall

³⁴We did have to free up the scale of marriage shocks in 1990-1996 in order to give the model a chance to match the data.

household inequality, while they fueled between-household inequality. Figure 9 (orange bars) shows that these effects are also quantitatively sizable. If only home production had changed over time, within-household inequality would have declined by 30% (accounting for more than the observed change) and the gender wage gap by 14% (accounting for more than 70% of the observed drop). Home production changes were thus the biggest driver of the decline in gender inequality. In turn, home production shifts put upward pressure on between-household inequality, accounting for almost 20% of the observed increase. But since this effect was dominated by the downward pressure on within-inequality, the net effect of technological change in home production on overall household income inequality was negative. Splitting home production further into the contributions of our model’s key parameters θ and ρ (Figure 13 in Appendix D.2) reveals that changes in relative productivity parameter θ were the main driver of the inequality shifts (accounting for around 2/3 of the total home production effects), while the impact of complementarity parameter ρ was smaller but still sizable (around 1/3 of the effects). The effects of changes in the labor market wedge ψ on inequality (yellow bars, Figure 9)—while qualitatively similar to those of home production technology—were quantitatively smaller. Finally, changes in labor market technology (blue bars), especially increase in γ_1 and A_z , fueled inequality across the board, significantly pushing up household income variance (through both the between and within components) and preventing gender inequality from falling even further. Thus, technological change in home production and in labor market production have pushed inequality, and especially gender disparities, in opposite directions. Our comparative statics in Section 8.1.2 clarify the mechanism regarding why the estimated changes

Figure 9: Mechanism behind Inequality Changes



of home production technology and the labor wedge push toward more gender equality. Both changes induced women to work more (leading to a decline in the gender gap of labor hours), which in turn caused women to sort relatively better in the labor market (thus reducing the gender gap in labor sorting). More gender parity in labor market outcomes, in turn, strengthened the desire for positive sorting in marriage, reinforcing the push toward more equal labor (and home) hours across gender. Figure 8,

right panel, demonstrates that these shifts were not only present in the model (purple bars) but also in the data (turquoise bars). Our evidence and estimates suggest that Germany underwent significant changes over the last decades toward an equilibrium that resembles the monotone equilibrium from our theory, with stronger home production complementarities and, consequently, increased marriage sorting as well as stronger comovements of spouses’ hours, labor market sorting, and wages.

We end by returning to a core feature of our model: equilibrium sorting in both labor and marriage markets. We assess the quantitative role of changes in marriage and labor market sorting for inequality shifts. Between 1990-1996 and 2010-2016, positive marriage sorting increased by around 10% and positive labor sorting by 8%. We compute the elasticity of each inequality outcome with respect to sorting in each market as $(\% \Delta \text{Inequality}) / (\% \Delta \text{Labor Sorting})$ and $(\% \Delta \text{Inequality}) / (\% \Delta \text{Marriage Sorting})$. *Inequality* refers to one of our four inequality outcomes (gender wage gap, household income variance, within/between component) and the percentage change is computed between the baseline model in 2010-2016 and the counterfactual model. This counterfactual inputs the estimated parameters from 2010-2016 but keeps either labor market or marriage market sorting constant at the past period’s (1990-1996) level.³⁵ In this way, we isolate the role of the observed changes in sorting for inequality shifts.

Table 4: Elasticity of Inequality with Respect to Sorting

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

Table 4 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 overall inequality and between-household inequality. For instance, a 1% increase in marriage sorting has decreased within-household inequality by 0.117%, while it increased between-household inequality by 0.123%. The elasticity of the gender wage gap is also negative, albeit smaller. Stronger marriage market sorting generated more balanced labor market outcomes—in hours, sorting, and pay—across 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 0.735%. In turn, a 1 % increase in labor sorting reduced the gender wage gap by 1.235% and within inequality by 0.736%. Surprisingly at first sight, the increase in labor sorting over the past decades significantly *narrowed* gender disparities. The reason is that this increase was predominantly driven by women’s improved labor sorting (i.e., the gender gap in labor sorting has declined over time; Figure 8, right panel), helping them to catch up with men’s pay. Stronger positive sorting between workers and jobs—when over-proportionally benefitting women—can spur gender convergence in labor market outcomes.

³⁵To implement the past marriage sorting in the counterfactual model, we adjust σ_{β}^M . In turn, to implement past labor sorting, we take into account labor market matching $\mu(\bar{s})$ from the past period when computing the wage function.

9 Conclusion

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

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

We then examine the nature of home production in the data. To this end, we estimate our model on data from today's Germany and find that spouses' home hours are strategic *complements*, pushing toward positive sorting in both markets and comovement of labor hours of spouses. 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 among partners. Home production complementarities induce spouses to split their time similarly between work in the market and at home; they also increase marriage sorting and reduce the gender gap in labor sorting, both of which further mitigate gender disparities.

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

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 amplify or mitigate the effects of aggregate labor market shocks and impact risk sharing within the household—issues we plan to address in future work.

References

- D. Aaronson and E. French. The effect of part-time work on wages: Evidence from the social security rules. *Journal of Labor Economics*, 22(2):329–252, 2004.
- J. Adda, C. Dustmann, and K. Stevens. The career costs of children. *Journal of Political Economy*, 125(2):293–337, 2017.
- S. Albanesi and C. Olivetti. Home production, market production and the gender wage gap: Incentives and expectations. *Review of Economic Dynamics*, 12(1):80 – 107, 2009.
- I. Andrews, M. Gentzkow, and J. M. Shapiro. Measuring the sensitivity of parameter estimates to estimation moments. *The Quarterly Journal of Economics*, 132(4):1553–1592, 2017.
- N. Angelov, P. Johansson, and E. Lindahl. Parenthood and the gender gap in pay. *Journal of Labor Economics*, 34(3):545–579, 2016.
- G. Anyaegbu. Using the OECD equivalence scale in taxes and benefits analysis. *The Labour Gazette*, 4(1):49–54, 2010.
- J. Bagger and R. Lentz. An empirical model of wage dispersion with sorting. *The Review of Economic Studies*, 86(1):153–190, 2018.
- G. S. Becker. A theory of marriage: Part I. *Journal of Political Economy*, 81(4):813–846, 1973.
- M. Bertrand, C. Goldin, and L. F. Katz. Dynamics of the gender gap for young professionals in the financial and corporate sectors. *American Economic Journal: Applied Economics*, 2(3):228–55, 2010.
- A. Bick, A. Blandin, and R. Rogerson. Hours and wages. 2020.
- S. Bonhomme, T. Lamadon, and E. Manresa. A distributional framework for matched employer employee data. *Econometrica*, 87(3):699–739, 2019.
- J. N. Brown and H. S. Rosen. On the Estimation of Structural Hedonic Price Models. *Econometrica*, 50(3):765–68, May 1982.
- M. Browning, P.-A. Chiappori, and Y. Weiss. *Economics of the Family*. Cambridge Books. Cambridge University Press, 2014.
- D. Card, J. Heining, and P. Kline. Workplace heterogeneity and the rise of west german wage inequality. *The Quarterly Journal of Economics*, 128(3):967–1015, 2013.
- H. Chade, J. Eeckhout, and L. Smith. Sorting through search and matching models in economics. *Journal of Economic Literature*, 55(2):493–544, June 2017.
- P. Chiappori, M. Iyigun, and Y. Weiss. Investment in schooling and the marriage market. *The American Economic Review*, 99(5):1689–1713, 2009.
- P.-A. Chiappori, B. Salanié, and Y. Weiss. Partner choice, investment in children, and the marital college premium. *American Economic Review*, 107(8):2109–67, August 2017.
- P.-A. Chiappori, M. Costa-Dias, and C. Meghir. The marriage market, labor supply, and education choice. *Journal of Political Economy*, 126(S1):S26–S72, 2018.

- P. A. Chiappori, M. C. Dias, and C. Meghir. Changes in assortative matching: Theory and evidence for the US and the UK. 2020.
- H. L. Cole, G. J. Mailath, and A. Postlewaite. Efficient non-contractible investments in large economies. *Journal of Economic Theory*, 101(2):333 – 373, 2001.
- P. Cortés and J. Pan. When time binds: Substitutes for household production, returns to working long hours, and the skilled gender wage gap. *Journal of Labor Economics*, 37(2):351–398, 2019.
- G. Cubas, C. Juhn, and P. Silos. Coordinated work schedules and the gender wage gap. NBER Working Paper 26548, 2019.
- M. C. Dias, R. Joyce, and F. Parodi. The gender pay gap in the UK: Children and experience in work. IFS Working Papers W18/02, 2018.
- 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. Working Paper 23636, National Bureau of Economic Research, 2017.
- R. Fernández and R. Rogerson. Sorting and Long-Run Inequality*. *The Quarterly Journal of Economics*, 116(4):1305–1341, 11 2001.
- R. Fernández, N. Guner, and J. Knowles. Love and money: A theoretical and empirical analysis of household sorting and inequality. *The Quarterly Journal of Economics*, 120(1):273–344, 2005.
- L. Flabbi, C. J. Flinn, and M. Salazar-Saenz. Simultaneous search in the labor and marriage markets with endogenous schooling decisions. 2020.
- G.-L. Gayle and L. Golan. Estimating a Dynamic Adverse-Selection Model: Labour-Force Experience and the Changing Gender Earnings Gap 1968–1997. *The Review of Economic Studies*, 79(1):227–267, 2011.
- G.-L. Gayle and A. Shephard. Optimal taxation, marriage, home production, and family labor supply. *Econometrica*, 87(1):291–326, 2019.
- D. Gicheva. Working long hours and early career outcomes in the high-end labor market. *Journal of Labor Economics*, 31(4):785–824, 2013.
- C. Goldin. A grand gender convergence: Its last chapter. *American Economic Review*, 104(4):1091–1119, April 2014.
- M. Goussé, N. Jacquemet, and J.-M. Robin. Household labour supply and the marriage market in the UK, 1991–2008. *Labour Economics*, 46:131 – 149, 2017a.
- M. Goussé, N. Jacquemet, and J.-M. Robin. Marriage, labor supply, and home production. *Econometrica*, 85(6):1873–1919, 2017b.

- J. Greenwood. *Evolving Households The Imprint of Technology on Life*. MIT Press, 2019.
- J. Greenwood, N. Guner, G. Kocharkov, and C. Santos. Technology and the changing family: A unified model of marriage, divorce, educational attainment, and married female labor-force participation. *American Economic Journal: Macroeconomics*, 8(1):1–41, 2016.
- J. Greenwood, N. Guner, and G. Vandenbroucke. Family economics writ large. *Journal of Economic Literature*, 55(4):1346–1434, December 2017.
- M. Hagedorn, T. H. Law, and I. Manovskii. Identifying equilibrium models of labor market sorting. *Econometrica*, 85(1):29–65, 2017.
- J. J. Heckman. Sample selection bias as a specification error. *Econometrica*, 47(1):153–161, 1979.
- G. L. Hunt. Alternative nested logit model structures and the special case of partial degeneracy. *Journal of Regional Science*, 40(1):89–113, 2000.
- M. Keane, P. Todd, and K. I. Wolpin. The structural estimation of behavioral models: Discrete choice dynamic programming methods and applications. volume 4A, chapter 04, pages 331–461. Elsevier, 1 edition, 2011.
- H. Kleven, C. Landais, and J. E. Søgaaard. Children and gender inequality: Evidence from Denmark. *American Economic Journal: Applied Economics*, 11(4), 2019.
- P. Legros and A. F. Newman. Beauty is a beast, frog is a prince: Assortative matching with nontransferabilities. *Econometrica*, 75(4):1073–1102, 07 2007.
- I. Lindenlaub. Sorting multidimensional types: Theory and application. *The Review of Economic Studies*, 84(2):718–789, 2017.
- I. Lindenlaub and F. Postel-Vinay. Multidimensional sorting under random search. 2020.
- J. Lise and F. Postel-Vinay. Multidimensional skills, sorting, and human capital accumulation. *American Economic Review*, Forthcoming.
- J. Lise and S. Seitz. Consumption inequality and intra-household allocations. *The Review of Economic Studies*, 78(1):328–355, 01 2011.
- S. Lundberg, R. Pollak, and J. Stearns. Family inequality: Diverging patterns in marriage, cohabitation, and childbearing. 2016.
- A. Mas and A. Pallais. Valuing alternative work arrangements. *American Economic Review*, 107(12): 3722–59, 2017.
- M. Mazzocco. Household intertemporal behaviour: A collective characterization and a test of commitment. *The Review of Economic Studies*, 74(3):857–895, 2007.
- D. McFadden. A method of simulated moments for estimation of discrete response models without numerical integration. *Econometrica*, 57(5):995–1026, 1989.
- 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. Working Paper, 2021.
- OECD. Starting strong IV. Early childhood education and care data country note. Germany. Technical report, OECD Publishing, Paris, 2016.
- A. Pakes and D. Pollard. Simulation and the asymptotics of optimization estimators. *Econometrica*, 57(5):1027–1057, 1989.
- L. Pilossoph and S. L. Wee. Assortative matching and income inequality: A structural approach. 2019a.
- L. Pilossoph and S. L. Wee. Household search and the marital wage premium. 2019b.
- G. Ramey and V. Ramey. The rug rat race. *Brookings Papers on Economic Activity*, 41:129–199, 08 2009.
- A. Reynoso. The impact of divorce laws on the equilibrium in the marriage market. 2019.
- S. Rosen. Hedonic prices and implicit markets: Product differentiation in pure competition. *Journal of Political Economy*, 82(1):34–55, Jan.-Feb. 1974.
- L. S. Shapley and M. Shubik. The assignment game I: The core. *International Journal of Game Theory*, 1(1):111–130, 1971.
- A. Voena. Yours, mine, and ours: Do divorce laws affect the intertemporal behavior of married couples? *American Economic Review*, 105(8):2295–2332, 2015.

Appendix

A Theory

A.1 Spouses' Hours as Strategic Complements or Substitutes

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

$$\begin{aligned}
 0 &= (w''e_h^2 + w'e_{hh})dh_f + 2p_{\ell_m\ell_f}dh_m + 2p_{\ell_f\ell_f}dh_f \\
 \Leftrightarrow \frac{dh_f}{dh_m} &= -\frac{2p_{\ell_m\ell_f}}{w''e_h^2 + w'e_{hh} + 2p_{\ell_f\ell_f}}
 \end{aligned} \tag{15}$$

$$\begin{aligned}
 0 &= (w''e_h^2 + w'e_{hh})dh_m + 2p_{\ell_m\ell_f}dh_f + 2p_{\ell_m\ell_m}dh_m \\
 \Leftrightarrow \frac{dh_m}{dh_f} &= -\frac{2p_{\ell_m\ell_f}}{w''e_h^2 + w'e_{hh} + 2p_{\ell_m\ell_m}}
 \end{aligned} \tag{16}$$

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

A.2 Proof of Proposition 1

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 .³⁶

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 couple's problem is

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

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 thus 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 couple'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

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

³⁶As 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 $w(\tilde{x}) = \int_0^{\tilde{x}} z_{\tilde{x}}$ 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 $\pi(y) = \int^y z_y(\mu^{-1}(t), t) dt$, 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').³⁷ 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 characterizes the optimal choices. It 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 strictly supermodular in (ℓ_m, ℓ_f)). So as above, SSCP holds and thus the unique solution (in any regular

³⁷SSCP 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 $w(e(x'_m, h'_m)) + V(x'_f, h'_m) - (w(e(x'_m, h''_m)) + V(x'_f, h''_m)) > 0$.

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 —then 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 also strictly increasing in h_m since p is strictly supermodular. Thus, h_f is strictly increasing in both x_f and h_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 $(\eta(x_f), x_f)$.

And, 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

A.3 Proof of Proposition 2

The proof is analogous to that of Proposition 1, which is why we shorten the argument and highlight the modifications.

LABOR MARKET PROPERTIES. The argument for positive labor market sorting follows that of Proposition 1 and is therefore omitted.

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 couple's problem is

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

where we replaced $\tilde{x}_m = e(x_m, h_m)$ and $\tilde{x}_f = e(x_f, h_f)$, and where the value of this problem is denoted by $\Phi(x_m, x_f)$. If hours are strictly *decreasing* in the partner's type, then $\Phi(x_m, x_f)$ is strictly submodular and thus *negative* sorting emerges in the marriage market.

PROPERTIES OF THE HOURS FUNCTIONS. We will show that h_m strictly increases in x_m and strictly *decreases* in x_f . To do so, write the couple'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 again 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

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

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 . The objective function is additively separable in (x_m, h_m) and (x_f, h_m) . 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, implying that h_m is increasing in x_m . And if V is strictly *submodular* in (x_f, h_m) , it will follow that h_m is decreasing in x_f (we will verify the submodularity of V below).

To show that h_m is *strictly* increasing in x_m , we follow the same argument as in the Proposition 1. And to show that h_m is *strictly* decreasing in x_f , we also follow this line of argument, only taking into account a submodular V (instead of supermodular V).

Let us now consider the inner maximization problem:

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

Note that the first term of the objective function is again strictly supermodular in (x_f, h_f) (since $w(e(\cdot, \cdot))$ is the composition of a supermodular and a convex function); and the second one is strictly submodular in (h_m, h_f) (under our assumption that p is strictly submodular in (ℓ_m, ℓ_f)). So, any optimal h_f in this problem is increasing in x_f and decreasing in h_m . And since h_m is strictly increasing in x_m (see above), we obtain that h_f is decreasing 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 —then it must be strictly increasing in x_f and strictly decreasing in h_m . ‘Strictly increasing’ follows from the same argument as in the proof of Proposition 1. In turn, ‘strictly decreasing’ follows from the same line of argument, only taking into account that p is strictly submodular. Thus, h_f is strictly increasing in x_f and strictly decreasing in h_m .

To complete the proof, it remains to show that V is strictly submodular 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 (Milgrom and Segal, 2002, Corollary 4(iii)), 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 submodular and h_f strictly increasing in x_f , it follows that V is strictly submodular. Hence, the premise made above when analyzing the choice of h_m in the outer maximization problem holds.

The remaining part of the proof is analogue to that of Proposition 1: The shown hours properties hold also along the equilibrium assignment; and since hours are strictly monotonic in each attribute, it follows that \tilde{N} is indeed atomless. \square

B Solution of the Quantitative Model

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

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

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

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

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

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

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

B.2 Optimal Household Choices ([hh])

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

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

When solving the household problem, agents take as given wage function w , the marriage market matching distribution η , and the marriage market clearing price v . Couples solve problem (12) and singles solve problem (13). Replacing the constraints in the objective function and noting the transferable utility structure of the problem, the collective problem of couple (s_f, s_m) , after labor supply shocks realize, is given by:

$$\max_{h_m, h_f} w(\tilde{s}_m) + w(\tilde{s}_f) + 2p^M(1 - h_m, 1 - h_f) + \delta^{h_m} + \delta^{h_f}, \quad (17)$$

where $w(\tilde{s}_m)$ and $w(\tilde{s}_f)$ depend on hours through the effective human capital types (14).

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

$$\max_{h_f} w(\tilde{s}_f) + p^U(1 - h_f) + \delta^{h_f}, \quad (18)$$

and the problem of a single man s_m is given by

$$\max_{h_m} w(\tilde{s}_m) + p^U(1 - h_m) + \delta^{h_m}. \quad (19)$$

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

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

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

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

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

$$\bar{u}_s^t(\mathbf{h}^t) = \begin{cases} w(\tilde{s}_i) + p^U(1 - h_i) & \text{if } t = U \\ w(\tilde{s}_m) + w(\tilde{s}_f) + 2p^M(1 - h_m, 1 - h_f) & \text{if } t = M. \end{cases} \quad (20)$$

We obtain the optimal private consumption and labor supply (c_m, c_f, h_m, h_f) for each household by solving problems (17)-(19). Given our assumption that the labor supply shock distribution is Type-I extreme value, we then obtain the probability that household type $t \in \{M, U\}$ with human capital type $\mathbf{s} \in \{\mathcal{S} \cup \emptyset\}^2$ chooses hours alternative $\mathbf{h} \in \{\mathcal{H} \cup \emptyset\}^2$ as:

$$\pi_s^t(\mathbf{h}^t) = \frac{\exp(\bar{u}_s^t(\mathbf{h}^t)/\sigma_\delta)}{\sum_{\tilde{\mathbf{h}}^t \in \{\mathcal{H} \cup \emptyset\}^2} \exp(\bar{u}_s^t(\tilde{\mathbf{h}}^t)/\sigma_\delta)}. \quad (21)$$

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

$$\eta_s \times \pi_s^t(\mathbf{h}^t).$$

From this distribution of *household* labor supply we back out the distribution of *individual* labor supply. To do so, we compute the fraction of men and women of each *individual* human capital type, s_i in household \mathbf{s} , optimally choosing each individual hours alternative h_i that is associated with household labor supply \mathbf{h} . Given the distribution of individual labor supply, we can compute the distribution of effective human capital types, \tilde{N}_s . First, note that the support of the distribution is obtained by applying functional forms (14) for any combination of *individual* hours and skill types. Second, to each point in the support of \tilde{s} , we attach the corresponding frequencies from the individual labor supply distribution backed out as explained above.

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

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

In the marriage stage, individuals draw idiosyncratic taste shocks for partners and singlehood, β_i^s , with $i \in \{f, m\}$ and $s \in \{\mathcal{S} \cup \emptyset\}$. At this stage, labor supply shocks are not yet realized. As a result, the *ex ante* economic value from a marriage of types (s_f, s_m) is the expected value of (17); and the ex-ante economic value from female and male singlehood is the expected value of (18) and (19). In both cases, the expectation is taken over the distribution of δ -shocks. Denoting the utility transfer to a female spouse of type s_f by $v(s_f)$, the values of being married (economic plus non-economic) for a female type s_f and a male type s_m in couple (s_f, s_m) are given by

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

where $\kappa = 0.57722$ is the Euler constant, \bar{u} is defined in (20) and \mathbb{E}_δ indicates that the expectation is taken over the distribution of δ -shocks.

In turn, the values of being single for woman s_f and man s_m are respectively given by:

$$\begin{aligned}\Phi_f(\emptyset, s_f) + \beta_f^\emptyset &:= \sigma_\delta \left[\kappa + \log \left(\sum_{\mathbf{h}^U \in \mathcal{H}} \exp\{\bar{u}_{(s_f, \emptyset)}^U(\mathbf{h}^U)/\sigma_\delta\} \right) \right] + \beta_f^\emptyset \\ \Phi_m(s_m, \emptyset) + \beta_m^\emptyset &:= \sigma_\delta \left[\kappa + \log \left(\sum_{\mathbf{h}^U \in \mathcal{H}} \exp\{\bar{u}_{(\emptyset, s_m)}^U(\mathbf{h}^U)/\sigma_\delta\} \right) \right] + \beta_m^\emptyset.\end{aligned}$$

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

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

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

$$\begin{aligned}\max \quad & \sum_{(s, s') \in [0, 1]^2} \eta_{(s, s')} \times (\Phi(s, s') + \tilde{\beta}) \\ s.t. \quad & \sum_{s \in \mathcal{S}} \eta_s = 1/2 \\ & \sum_{s' \in \mathcal{S}} \eta_{s'} = 1/2,\end{aligned}$$

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

We obtain the equilibrium matching in the marriage market, η , by solving this linear program, taking prices and allocations in households and the labor market, $(w, \mu, h_f, h_m, \tilde{N}_s)$, as given.

B.4 General Equilibrium of the Model

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

B.4.1 Trembling Effective Types

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

To address this issue, we develop a trembling strategy. We draw a small random sample of women and men and force them to supply a suboptimal amount of hours from the set of unchosen hours in each iteration. In practice, for each group of women with skill type s_f and each group of men with skill type s_m , we track their optimal choices for a given wage function and determine the hours that were *not*

chosen with positive probability. We then draw a 1% random sample of women and men within each of those skill types (the ‘tremblers’) and assign them uniformly to the unchosen hours, forcing them to choose sub-optimal hours. Finally, we construct the distribution of effective types \tilde{N}_s by taking into account both ‘trembling’ effective types and ‘optimal’ effective types.

B.4.2 Fixed Point Algorithm

To solve for the general equilibrium, we denote by \tilde{N}_s^* the distribution of *realized* effective types (based on *optimal* hours choices, as opposed to *trembling* hours choices). Similarly, we denote by w^* the wage as a function of *realized* effective types only, while the *full support* wage function is denoted by w . The fixed point algorithm we designed to solve for the equilibrium is as follows:

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

At any round $r \geq 1$

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

C Estimation

C.1 Identification

C.1.1 Identification of the Worker and Job Distribution

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

C.1.2 Proof of Proposition 3

As is standard in discrete choice models, our argument relies on variation in choice probabilities of hours and partners. That is, generically, individuals of a given type choose different hours alternatives with different probabilities; and individuals of a given type choose different marriage partners with different probabilities and, moreover, these choice probabilities also vary across types.

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

$$w'(\tilde{s}) = z_{\tilde{x}}(\tilde{s}, \mu(\tilde{s})) \quad (22)$$

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 (22) after applying a log transformation and taking into account measurement error:

$$\log(\hat{w}_{\tilde{s}}(\tilde{s})) = \log(z_{\tilde{x}}(\tilde{s}, \mu(\tilde{s}))) + \epsilon \quad (23)$$

where, for concreteness, we assume the functional form $z(\tilde{s}, y) = A_z \tilde{s}^{\gamma_1} y^{\gamma_2} + K$ (see main text), and where we treat \tilde{s} and the matching μ as observed. Note that this functional form of z circumvents the identification problem of [Rosen \(1974\)](#), discussed in [Brown and Rosen \(1982\)](#) and [Ekeland et al. \(2004\)](#), since the slope of the wage gradient in \tilde{s} is *not* equal to the slope of the marginal product in \tilde{s} . We assume that ϵ is the measurement error of the marginal wage, with mean zero and uncorrelated with the right-hand-side (RHS) variables. Regression (23) 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 \tilde{g} , who—due to PAM in the labor market—matches with the lowest firm type $y = 0$ and thus $z(\tilde{g}, 0) = K$, meaning any positive wage $w(\tilde{g}) > 0$ can only be attributed to K , i.e., $w(\tilde{g}) = K$.

Identification of the Female Productivity Wedge. We can identify ψ from the within (s, h) -type (agents with the same s and same work hours h) wage gap across gender. Denote the gender wage gap within individuals of hours-human-capital type $(s, h) = (\hat{s}, \hat{h})$ 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 which worker matches to which firm. 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:

$$w(\psi\hat{s}\hat{h}) = w(\hat{s}\hat{h})(1 - gap(\hat{s}, \hat{h})) \quad (24)$$

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: either there is an interior gap, $0 < gap(\hat{s}, \hat{h}) < 1 - K/w(\hat{s}\hat{h})$, and so by the Intermediate Value Theorem there exists a unique $\psi \in (0, 1)$ for which (24) holds; or, the minimal gap $gap(\hat{s}, \hat{h}) = 0$ pins down $\psi = 1$; or the maximal gap $gap(\hat{s}, \hat{h}) = 1 - K/w(\hat{s}\hat{h})$ pins down $\psi = 0$. Thus, ψ is identified from gender wage gaps of agents with the same (s, h) -combination.

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

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

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

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

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

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

$$\pi_s^t(\mathbf{h}^t) = \frac{\exp(\bar{u}_s^t(\mathbf{h}^t)/\sigma_\delta)}{\sum_{\tilde{\mathbf{h}}^t \in \{\mathcal{H} \cup \emptyset\}^2} \exp(\bar{u}_s^t(\tilde{\mathbf{h}}^t)/\sigma_\delta)} \quad (25)$$

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

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

$$\begin{aligned}
\frac{\pi_{\mathbf{s}}^U(\mathbf{h}^U)}{\pi_{\mathbf{s}}^U(\mathbf{0})} &= \frac{\exp(\bar{u}_{\mathbf{s}}^U(\mathbf{h}^U)/\sigma_{\delta})}{\exp(\bar{u}_{\mathbf{s}}^U(\mathbf{0})/\sigma_{\delta})} \\
\log\left(\frac{\pi_{\mathbf{s}}^U(\mathbf{h}^U)}{\pi_{\mathbf{s}}^U(\mathbf{0})}\right) &= \frac{\bar{u}_{\mathbf{s}}^U(\mathbf{h}^U) - \bar{u}_{\mathbf{s}}^U(\mathbf{0})}{\sigma_{\delta}} \\
&= \frac{w(s_m h_m) - w(s_m 0) + p^U(1 - h_m) - p^U(1 - 0)}{\sigma_{\delta}} \\
&= \frac{w(s_m h_m) + p^U(1 - h_m) - p^U(1)}{\sigma_{\delta}}
\end{aligned} \tag{26}$$

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

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

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 (it is the effect of men's human capital on wages given the hours choice $\mathbf{h}^U \neq \mathbf{0}$) is also observed and different from zero as the wage strictly increases in human capital. Thus, σ_{δ} is identified.

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

$$\begin{aligned}
\frac{\pi_{\mathbf{s}}^M(\mathbf{h}^M)}{\pi_{\mathbf{s}}^M(\mathbf{1})} &= \frac{\exp(\bar{u}_{\mathbf{s}}^M(\mathbf{h}^M)/\sigma_{\delta})}{\exp(\bar{u}_{\mathbf{s}}^M(\mathbf{1})/\sigma_{\delta})} \\
\log\left(\frac{\pi_{\mathbf{s}}^M(\mathbf{h}^M)}{\pi_{\mathbf{s}}^M(\mathbf{1})}\right) &= \frac{w(\psi s_f h_f) - w(\psi s_f) + w(s_m h_m) - w(s_m) + 2p^M(1 - h_m, 1 - h_f)}{\sigma_{\delta}}
\end{aligned} \tag{27}$$

where we used that $2p^M(0, 0) = 0$ by assumption in our quantitative model. Note that the LHS of (27) (relative choice probabilities) is observed, and on the RHS, wages of men and women with types (s_f, s_m) conditional on hours are also observed in the data, and σ_{δ} is known at this stage. Thus, home production function p^M is non-parametrically identified since we can specify (27) for all hours alternatives \mathbf{h}^M chosen in the data. Note that we can identify p^M from a couple of any type $\mathbf{s} = (s_f, s_m)$.

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

Identification of the Scale of the Marriage Taste Shock. We now show that σ_β^M is identified once the parameters of the utilities are identified.

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

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

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

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

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

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

Then, we can compare the ratio of these choice probabilities across male types (s'_m, s''_m):

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

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

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

The LHS is observed. Moreover, all objects on the RHS are either observed (wages) or identified at

this stage (home production function and σ_δ), except σ_β^M . We can solve this equation for σ_β^M , giving a unique solution. Thus, σ_β^M is identified.

C.2 Estimation of Worker Types

Sample Selection. Our sample consists of individuals in the GSOEP from 1984-2018 who are between 20 and 60 years old and are either married/cohabiting or single. We exclude individual-year observations when the individual indicated self-employment and when they worked in poorly defined occupations ($\text{k1db92} \geq 9711$). We also exclude observations with missing information on education or with missing (not zero) labor force experience. Our panel consists of around 212,000 person-year observations.

Method. We aim to estimate the distribution of workers’ human capital N_s . To do so, we need to estimate the returns to education to obtain a proxy for x in our model, as well as individual unobserved heterogeneity, capturing ν in our model. We use our theoretical wage function as guidance for specifying the empirical log hourly wage of individual i at time t (where t is a year in our sample) as

$$\ln w_{it} = \nu_i + \sum_{ed \in \{voc, c\}} \alpha^{ed} x_{it}^{ed} + \beta_1 h_{it} + \beta_2 h_{it}^2 + \beta'_z Z_{it} + \kappa_s + \rho_t + \epsilon_{it}, \quad (28)$$

where x_{it}^{ed} are indicator variables for the education group of an individual (meant to capture x in our model) at time t . Coefficient α^{ed} gives the ‘value’ of education ed in terms of log wage units, where $0 < \alpha^{voc} < \alpha^c$ would indicate positive returns to education. While these coefficients indicate the average return to education for all individuals in a certain education group, ν_i is a person fixed effect that captures unobserved time-invariant ability, with model counterpart ν . In turn, h_{it} denotes weekly labor hours (which captures the time ‘investment’ in productivity in our model). Finally, Z_{it} are time-varying controls for the individual, κ_s and ρ_t are state and time fixed effects, and ϵ_{it} is a mean-zero error term.³⁸ We implement (28) as a panel wage regression and deal with three challenges below: selection into labor force participation; endogeneity of hours worked; agents who hold jobs for less than two periods in our sample.

Key Variables. For *weekly hours*, we use reported actual hours which, when positive, we winsorized by 10 hours from below and 60 hours from above to deal with outliers. For *labor force experience* we use the reported labor force experience, and we impute it by potential experience if this information is missing.³⁹ For *education*, we use three categories: In the group of *low education*, there are those whose highest degree is lower secondary, high school or vocational with weakly less than 11 years of education (around 35%). In the group of *medium education* (which we sometimes denote by *voc*), there are those with vocational degree and above 11 years of education (around 44%). In the group of *high*

³⁸We do not include occupation fixed effects since in our model, conditional on \tilde{s} (which we control for here by controlling for (x, ν, h)), the wage does not depend on occupation in our competitive equilibrium. But even doing so—which we have done for robustness—does not significantly change the impact of x or ν on the hourly wage.

³⁹For men, potential and actual experience are almost perfectly correlated, which is why this imputation should work well for them. For women, the correlation is much lower, which is why we do not impute here.

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

Selection Equation. To account for selection into labor force participation in the wage regression, we first run a selection regression. To do so, we need an instrument that affects participation but is excluded from wage regression (28). Since the variation in participation in our sample is mainly due to women, we use the ‘progressiveness’ in an individual’s narrowly defined demographic cell. We proxy progressiveness by the share of females working in a narrowly defined demographic cell. Our cells are defined by a combination of state, year, age and an indicator whether a child under the age of 3 is in the household.⁴⁰ When defining this variable for a particular individual, we employ the ‘leave-one-out’ method and do not count the individual’s labor force participation when computing this statistic. We further drop cells with less than five observations. We end up with around 2,500 cells with more than 5 observations in each. Note that we experimented with additional cell characteristics (education and country of origin) but there is a trade-off between number of observations by cell and making the cells more specific to the demographic groups. Defining the cells by these additional variables would imply to drop more than twice as many observations due to small cell size.

Our assumption on this IV (i.e., our progressiveness variable) is that the following exclusion restriction holds: ‘progressiveness of an individual’s demographic cell’ only affects wages through labor force participation but not in other ways.

We run the following probit selection regression:

$$emp_{it} = \alpha share_{j(i)t} + \sum_{ed \in \{voc, c\}} \alpha^{ed} x_{it}^{ed} + \beta'_z Z_{it} + \kappa_s + \rho_t + \epsilon_{it} \quad (29)$$

where the dependent variable is an indicator of whether individual i is employed at time t , $share_{j(i)t}$ is the progressiveness measure in the demographic cell j of individual i at time t (given by the share of women working in the cell, see description of this IV above), x_{it}^{ed} captures education indicators (‘low education’ is the reference group), Z_{it} is a vector of demographic individual controls (linear and quadratic labor force experience in years, household size) and κ_s and ρ_t are state and year fixed effects and ϵ_{it} is a mean-zero error term. We cluster standard errors on the cell level.

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

⁴⁰German children start kinder-garden when they are 3 years old. Before age 3, children are predominantly at home (in 2013, only 29% of children aged 0-2 were in daycare (OECD, 2016)), so age 3 is an important threshold when it comes to mothers’ labor force participation.

Table 5: Selection Regression

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

Notes: Results are based on regression (29). A *Cell* is defined by state, age, year and presence of children under 3 in the household. The regression includes year and state fixed effects. Standard errors clustered at cell level in parentheses. ***Significant at the 1% level.

Panel Wage Regression. One potential issue for our specification of the wage regression is that there may be confounding factors that impact both hours and wages. While we deal with time-invariant unobserved heterogeneity using the panel regression with individual fixed effects, time-varying unobserved heterogeneity—such as productivity shocks or health shocks—could still be problematic. To address this concern, we use an IV approach. In our model, there is a systematic relationship between the hours worked of an individual and the hours worked by their partner, so we use the partner’s hours as an instrument for own hours. Identification relies on changes in spousal labor hours over time. The identifying assumption is that conditional on the individual fixed effect and education, partner’s hours are exogenous in the wage regression and that partner’s labor hours impact own wage *only* through own labor hours, which is satisfied in our model. In particular, we instrument hours worked and hours worked squared by (i) the hours worked by the partner, (ii) the hours worked by the partner squared and (iii) whether a partner is present (where we set partner’s hours to zero in both cases, if partner is present but not working and if partner is not present). We therefore drop observations whose partner reports to be employed but has zero reported labor hours and observations whose partner has missing employment and hours information. Since we include individual fixed effects we also drop singleton observations (those who only show up in a single year of the panel). We further restrict the sample to those that are employed and have non-missing hourly wage. This leads to a sample of 133,214 person-year

observations. Based on our model wage function, we choose the following regression specification:

$$\ln w_{it} = \nu_i + \beta_0 IMR_{it} + \beta_1 h_{it} + \beta_2 h_{it}^2 + \sum_{ed \in \{voc, c\}} \alpha^{ed} x_{it}^{ed} + \beta'_z Z_{it} + \kappa_s + \rho_t + \epsilon_{it} \quad (30)$$

where ν_i is a individual fixed effect, IMR_{it} is the inverse mills ratio of individual i in year t from the selection probit regression above, h_{it} are weekly hours worked, and x_{it}^{ed} , Z_{it} , κ_s and ρ_t are as in the selection regression (29). Note that we could not include h_{it} into the selection equation since $h_{it} = 0$ versus $h_{it} > 0$ is a perfect predictor of employment. Nevertheless, based on our model, it is important to control for hours worked in the hourly wage regression. We again cluster standard errors on the cell level.

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

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

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

Table 6: Wage Regression

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

Notes: All specifications include year and state fixed effects. The table reports the IV version of regression (30), where we instrument Weekly Hours Worked and Weekly Hours Worked squared by the partner's Weekly Hours Worked (linear and squared) and whether the partner is present. Columns (1) and (2) are the first stage regressions while Column (3) is the second stage regression. Standard errors (clustered at the cell level) are in parentheses. A Cell is defined by state, year, age, and presence of children under 3 in the household. ***Significant at the 1% level. **Significant at the 5% level.

Table 7: Worker Distribution of s -Types by Education

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

Notes: *Low Educ* includes high school and vocational education with less than 11 years of schooling. *Medium Educ* is defined as vocational education with more than 11 years of schooling. *High Educ* is defined as college and more. We obtain six s -types by ordering their value of $s_i = \alpha^{ed} x_i^{ed} + \bar{v}_j^{ed}(i)$ where parameters α^{ed} are estimated based on regression (30) and where $\bar{v}_j^{ed}(i)$ is the fixed effect of individual i in education group ed and subgroup j , also computed from model (30) as the average individual fixed effects in that group.

C.3 Estimation of Job Types

Data and Sample Selection. Our main data source for measuring occupation types is the BIBB collected in 2012 by the German Federal Institute of Vocational Training, and the German Federal Institute for Occupational Safety and Health. This survey is representative of the German employed population. In particular, it contains data on task usage in 1,235 occupations defined by the 4-digit code

k1db92, which we also use for our analysis in the GSOEP. This data is reported by individuals who work in these occupations. In order to reduce the problem of noisy reporting, we drop occupations in which the task information is based on less than 5 individuals. We are left with task data for 613 occupations. These are the most common occupations and we will base our structural estimation exercise on the subset (608 occupations) that can be merged to the occupations held by individuals in our GSOEP sample, using the four-digit occupational codes k1db92 from the German Classification of Occupations 1992.

Task Data. The BIBB contains data on how intensely different 4-digit occupations use different types of tasks. These intensity measures are continuous and we normalize them to be on the unit interval. The reported tasks are comprised of: Detailed Work, Same Cycle, New Tasks, Improve Process, Produce Items, Tasks not Learned, Simultaneous Tasks, Consequence of Mistakes, Reach Limits, Work Quickly, Problem Solving, Difficult Decisions, Close Gaps of Knowledge, Responsibility for Others, Negotiate, Communicate.

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

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

Figure 10: Principal Component Analysis




Table 8: Correlation Across Alternative Measures

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

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

***Significant at the 1% level.

Table 9: Summary Statistics of Alternative Measures

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

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

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

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

We also considered determining the occupational types through an occupational fixed effect in the wage regression we used to recover worker types (equation (28)). This would have meant to run a two-way fixed effects regression. We chose our alternative approach that does not rely on occupational fixed effects for the following reasons: First, based on our model featuring a competitive labor market, the wage function does not depend on an occupational fixed effect/type when controlling for workers' effective types: all workers with the same effective type \tilde{s} should be matched to the same occupation. Second, the two-way fixed effects approach is known to be problematic under limited worker mobility (limited mobility bias) and when one is interested in sorting (since the correlation between worker and firm/occupation fixed effects is not accurately capturing sorting).

C.4 Internal Estimation

C.4.1 Parameters Set Outside the Model

Table 10: Exogenously Fixed Parameters

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

C.4.2 Results

Table 11: Targeted Moments

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

Notes: *LFP* stands for Labor Force Participation. Moments are computed as discussed in Section OD.1.1 and Table O4 of the Online Appendix. *Types* refer to human capital types, where each number 1-6 corresponds to the corresponding column of Table 7.

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

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

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

D Quantitative Analysis

D.1 Comparative Statics

THE EFFECT OF ψ . Last, we analyze the comparative statics of the female labor market wedge. Figure 11a in Appendix D.1 shows that eliminating the wedge (increasing ψ from our estimate $\psi = 0.84$ to $\psi = 1$) would reduce the gender gap by about 25%. There is a direct positive effect of ψ on female productivity—and thus wages—but also several indirect effects: First, the wife’s labor hours increase in productivity ψ relative to the husband’s, reducing the gender hours gap (panel e) and thus the gender wage gap. Second, the reduction in the gender hours gap leads to a decline in the labor market sorting gap (panel f), further curbing the gender wage gap. Third, smaller gender disparities in the labor market are associated with an increase in marriage market sorting (panel d), since in a world in which men and women are more equal, the motive for positive sorting strengthens. The increase in marriage sorting reinforces the drop in both hours and labor sorting gaps, further dampening the gender wage gap.

In Figures 11b and c, we study the effects of ψ on the variance of income, both within and across households. An increase in ψ leads to lower within-household inequality, but higher between-household inequality (driven by the increase in marriage sorting), with ambiguous effects on overall inequality.

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





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



D.2 Inequality Over Time

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

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

Notes: *LFP* stands for Labor Force Participation. Moments are computed as discussed Section OD.1.1 and Table O4 of the Online Appendix. The last column of the table reports the p-value of the hypothesis test of the differences between the data moments in the two samples being zero. We use a standard T-test for differences in means (M6), a standard Levene test for differences in variances (M7) and standard tests for differences in proportions (M11). We use a Fisher transformation to construct the test statistic for the differences in correlations between samples (M5 and M10). We use a two-sample Wald test for differences in ratios across samples (M1, M2, M3, M4, M8, M9, M12 and M13). To construct the statistic for the Wald tests for difference in ratios, we use bootstrap techniques for the variance estimation.

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

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

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

Figure 13: Inequality Changes Over Time: Detailed Decomposition

