Marriage Market and Labor Market Sorting: Online Appendix

Paula Calvo Ilse Lindenlaub Ana Reynoso

October 15, 2022

Throughout this Online Appendix, we indicate figures, tables, and equations within this appendix by O.# ('O' for Online). In turn, figures, tables, and equations from the main paper are denoted by just 1,2,.... Figures, tables, and equations from the main appendix are denoted by A.#.

OA Empirical Evidence

OA.1 Time Allocation by Marital Status and Education

Figure O.1: Weekly Mean Hours by Gender, Education and Marital Status

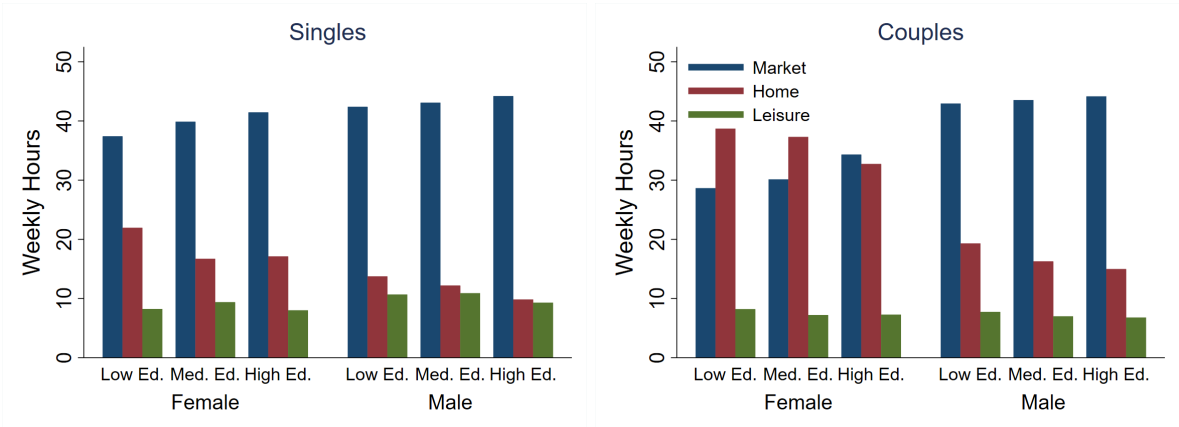

OA.2 Time Allocation and Wages

The left panel of Figure O.2 shows the wage penalties of various groups *relative* to men working full-time (the blue bar, whose height equals one), where full-time is defined as working between 37.5 and 60 hours per week. The height of the remaining bars in Figure O.2 (left) reflects the estimates of the regression of the logarithm of hourly wages on an indicator for part-time work, a gender indicator and the interaction between both. We control for age, marital status, birth place and education. While full-time women

have a wage penalty of 14.7 percentage points relative to full-time men, women who work part-time have an even larger penalty of 24 percentage points. Figure O.2 (right) shows that while less than 10% of employed men work part-time, more than 50% of employed women do so.

Figure O.2: Part-Time Wage Penalties Relative to Men Who Work full-time (left); Share of Workers in Full-/Part-Time Work (right)

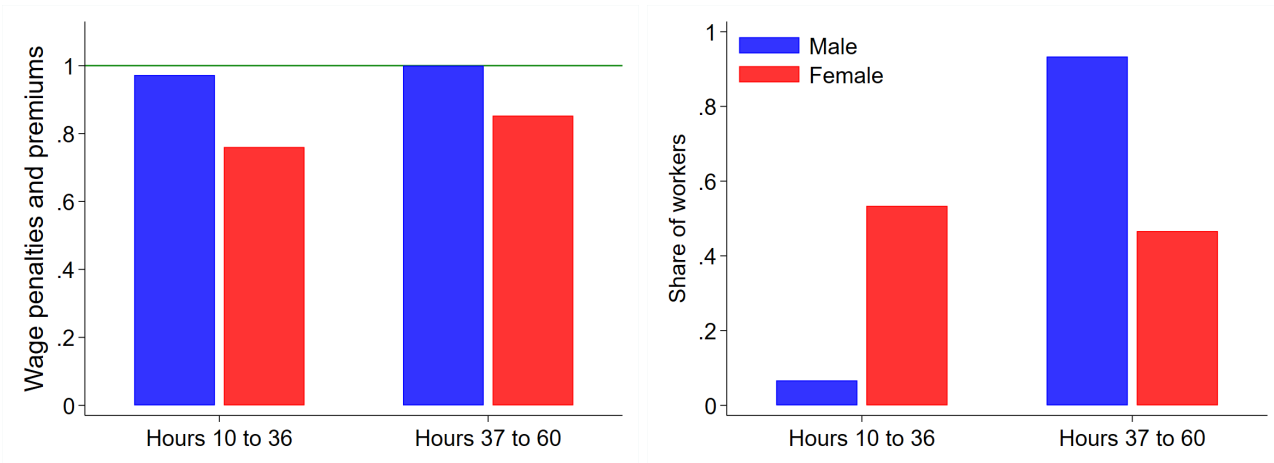

OB Theory

OB.1 TU Representation and Monotone Equilibrium

TU Representation and the Generalized Quasi-Linear Class. A broad class of utility functions makes our model TU-representable, according to which households' overall split between public and private consumption is independent of how private consumption is shared. It is well-known (Chiappori and Gugl, 2020) that the class of Generalized Quasi-Linear (GQL) utility functions is necessary and sufficient for the TU representation of the type of problem we are dealing with. The GQL is given by

$$u(c,p) = F(\alpha(p)c + \beta(p)),$$

where F, α, β are strictly increasing and β is strictly concave. Some concrete examples that fall into the GQL class and thus yield the TU representation in our model are:

$$u(c,p) = F(c+p), \quad \text{for any } F' > 0 \tag{O.1}$$

$$u(c,p) = \log(cp) \tag{O.2}$$

$$u(c,p) = \alpha(p)m(c) \quad \text{for any } m \text{ s.t. } m(xy) = m(x)m(y)$$

$$u(c,p) = \frac{c^{1-\gamma}}{1-\gamma}\frac{p^{1-\theta}}{1-\theta} = \frac{1}{(1-\gamma)(1-\theta)}\left(cp^{\frac{1-\theta}{1-\gamma}}\right)^{1-\gamma}, \gamma \geq 0, \theta \geq 0.$$

This highlights that many common utility functions lead to the TU representation in our model, including those that allow for flexible degrees of risk aversion.

Monotone Equilibrium Beyond the Quasi-Linear Class. The reason for why we focus on quasi-linear utility (O.1) within the GQL class is twofold: First, this assumption increases tractability because it guarantees the TU property. Second, and most importantly, it also allows us to obtain clean analytical (sufficient) conditions for monotone equilibrium, which captures well the features of the data we document in Section 2 of the paper. For several other functional forms in the GQL class, the conditions for monotone equilibrium would be more involved because they encapsulate forces toward household specialization and therefore toward the *non*-monotone equilibrium of our model.

For instance, consider log-utility (O.2). The complication with log-utility is that private and public consumption goods are *complements* in the utility function, generating a force towards household specialization (one partner works a lot at home so that p is high and the other one works a lot in the market so that c is also high), which pushes *against* the monotone equilibrium that—by definition—has the requirement that spouses' hours are complements and therefore co-move.

Nevertheless, more general utility functions from the GQL class (such as log-utility) do lead to predictions consistent with monotone equilibrium (and thus with the data) when considering a natural extension of our framework. This extension is to allow couples to purchase some of the public good in the market using their wages. We now denote by $\tilde{p} = p + c_p$ the public good, where c_p is the public good expenditure in the market while p is home-produced, as before. The idea is that the consumption good that is purchased in the market can not only be used for private consumption but also as an input for the public good. We assume for illustration that the public good produced at home and the one purchased in the market are perfect substitutes. This gives us analytical tractability and is plausible given the high degree of substitutability between home production and market purchases found in the literature (e.g., Chang and Schorfheide, 2003 and Greenwood, Guner, and Vandenbroucke, 2017).

We now take a closer look at our model with log-utility under this extension: We start with the household decision problem. The collective decision problem of any potential couple (x_m, x_f) that forms in the marriage market is given by

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

where we normalized prices to 1, as is often assumed in this literature (see, e.g., many of the applications in Browning, Chiappori, and Weiss, 2014).

Plugging the constraints into the objective, and using a monotone transformation, gives

$$\begin{aligned} &\max_{c_p, h_m, h_f} \log\left(\left(w(\tilde{x}_m) + w(\tilde{x}_f) - c_p - \frac{\exp(\overline{v})}{\tilde{p}}\right)\tilde{p}\right) \\ \Leftrightarrow \quad &\max_{c_p, h_m, h_f} \big(w(\tilde{x}_m) + w(\tilde{x}_f) - c_p\big)\big(c_p + p(1 - h_m, 1 - h_f)\big). \end{aligned}$$

Taking the FOC w.r.t. c_p gives: $c_p = (w(\tilde{x}_m) + w(\tilde{x}_f) - p(1-h_m, 1-h_f))/2$. Plugging c_p back into the objective (and applying another monotone transformation) gives:

$$\max_{h_m, h_f} w(\tilde{x}_m) + w(\tilde{x}_f) + p(1-h_m, 1-h_f),$$

which closely resembles our objective function from the baseline model. Thus, in the household stage, the conditions for hours to be increasing in own and partner's type remain the same. In the marriage market stage, the conditions for positive sorting are also the same. The conditions on primitives for monotone equilibrium are therefore *identical* to those in the baseline model with quasi-linear preferences.

The intuition for this result is straightforward: Substitutability in the public good $\tilde{p}$ between the component produced at home using hours, p, and the one that is purchased with wages, c_p, *undoes* the complementarity between p and c_m—or between p and c_f—in the utility function. So even though under log-utility, private and public consumption (c_m and $\tilde{p}$) are complements (and thus wages and $\tilde{p}$ are complements, pushing towards non-monotone equilibrium), the substitutability between p and c_p in $\tilde{p}$ counteracts these forces. This re-introduces a substitutability between wages (that purchase c_p and also c_m) and p. Therefore, if p is complementary in spouses' home production time, then this setting gives rise to the monotone equilibrium, in which spouses' hours are positively correlated, as in the baseline model.

As a result, what our simple functional form (quasi-linearity) captures is in essence a more general utility function from the GQL class when allowing for substitutability in the public good between the component produced at home and the component purchased in the market.

OB.2 Monotone Equilibrium in the Quantitative Extension of the Model

We now show that the key properties of our baseline model (Propositions 1 and 3 of the paper) are preserved in the quantitative extension of our model, when adjusted for the stochastic nature of hours and matching. Since Propositions 1 and 3 are flip sides of each other, we here focus on generalizing Proposition 1. We will show that under similar conditions as in the baseline model, the properties of monotone equilibrium hold *on average*.

We will maintain the following assumptions (see also Section 4.1). Marriage taste shocks and labor supply shocks follow type-I extreme-value distributions:

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

where

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

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

Proposition O1. *If p is strictly supermodular in (ℓ_m, ℓ_f) and z are strictly supermodular in $(\tilde{s}, y)$ and convex $\tilde{s}$, then the properties of monotone equilibrium are satisfied on average, i.e.*

1. *labor market: there is positive sorting between effective human capital $\tilde{s}$ and firm types y;*
2. *households: on average, labor hours h_i are increasing in own type s_i and in partner's type s_j, $i \neq j, i, j \in \{f, m\}$;*
3. *marriage market: on average, there is positive sorting in the sense that higher s_f match with higher s_m.*

Proof.

Part 1. Recall the firm's problem

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

Based on well-known arguments, given that z is assumed to be strictly supermodular in $(\tilde{s}, y)$, the optimal matching satisfies positive sorting in $(\tilde{s}, y)$.

Part 2. We want to show that $\widehat{\Pi}_m(h_m|s_m)$ and $\widehat{\Pi}_f(h_f|s_m)$ are decreasing in s_m, where

$$\widehat{\Pi}_m(h_m|s_m) \equiv \sum_{s_f} \Pi_m(h_m|s_m, s_f)\eta(s_f|s_m)$$
$$\widehat{\Pi}_f(h_f|s_m) \equiv \sum_{s_f} \Pi_f(h_f|s_m, s_f)\eta(s_f|s_m)$$

are marginal cdfs of male and female hours conditional on male (own and partner) type, so that higher male s_m-types are associated with stochastically higher labor market hours both for them and their partner (the argument for hours being increasing in female types s_f is analogous and omitted). Note that $\eta(s_f|s_m)$ is the conditional marriage matching probability mass function (pmf), and $\Pi_m(h_m|s_m, s_f)$ and $\Pi_f(h_f|s_m, s_f)$ are the marginal cdfs of male and female hours, conditional on partners' types,

$$\Pi_m(h_m|s_m, s_f) \equiv \sum_{\tilde{h}_m < h_m} \sum_{\tilde{h}_f} \pi_{(s_f, s_m)}(\tilde{h}_f, \tilde{h}_m)$$
$$\Pi_f(h_f|s_m, s_f) \equiv \sum_{\tilde{h}_f < h_f} \sum_{\tilde{h}_m} \pi_{(s_f, s_m)}(\tilde{h}_f, \tilde{h}_m),$$

obtained from the joint pmf of spouses' hours, $\pi_{(s_f, s_m)}(h_f, h_m)$ (see (A.7) of Section C.2.2):

$$\pi_{(s_f, s_m)}(h_f, h_m) \quad = \quad \frac{\exp(\overline{u}_{\mathbf{s}}(\mathbf{h})/\sigma_\delta)}{\sum_{\tilde{\mathbf{h}} \in \{\mathcal{H} \cup \emptyset\}^2} \exp(\overline{u}_{\mathbf{s}}(\tilde{\mathbf{h}})/\sigma_\delta)}.$$

We dropped the household type superscript t to reduce notation (we exclusively focus on couples here). We will now derive the conditions under which $\widehat{\Pi}_m(h_m|s_m)$ is decreasing in s_m step-by-step.

As we are interested in the impact of s_m on $\widehat{\Pi}_m(h_m|s_m)$, we apply the discrete chain rule to obtain:

$$\Delta_{s_m}\widehat{\Pi}_m(h_m|s_m) = \sum_{s_f}\Pi_m(h_m|s_m,s_f)\Delta_{s_m}\eta(s_f|s_m) + \sum_{s_f}\Delta_{s_m}\Pi_m(h_m|s_m,s_f)\eta(s_f|s_m) \tag{O.3}$$

where we denote the discrete derivative of a function $f(n)$ by $\Delta_n f(n) = f(n+1) - f(n)$. We want to establish conditions under which $\Delta_{s_m}\widehat{\Pi}_f(h_f|s_m) \leq 0$ in (O.3).

We can further simplify this expression by applying summation by parts to the first term. To do so, let's index the different female types by k, so $s_{f_k}, k \in \{0, ..., n\}$. Then:

$$\begin{aligned}\sum_{s_f}\Pi_m(h_m|s_m,s_f)\Delta_{s_m}\eta(s_f|s_m) &= \sum_{k=0}^{n}\Pi_m(h_m|s_m,s_{f_k})\Delta_{s_m}\eta(s_{f_k}|s_m)\\ &= \Pi_m(h_m|s_m,s_{f_n})\sum_{k=0}^{n}\Delta_{s_m}\eta(s_{f_k}|s_m) - \sum_{j=0}^{n-1}\Delta_{s_{f_j}}\Pi_m(h_m|s_m,s_{f_j})\sum_{k=0}^{j}\Delta_{s_m}\eta(s_{f_k}|s_m)\end{aligned}$$

where in the first term, we have

$$\sum_{k=0}^{n}\Delta_{s_m}\eta(s_{f_k}|s_m) = \sum_{k=0}^{n}\eta(s_{f_k}|s_{m+1}) - \sum_{k=0}^{n}\eta(s_{f_k}|s_m) = 0$$

and thus the first term vanishes due to the standard property of cdfs. Therefore,

$$\sum_{s_f}\Pi_m(h_m|s_m,s_f)\Delta_{s_m}\eta(s_f|s_m) = -\sum_{j=0}^{n-1}\Delta_{s_{f_j}}\Pi_m(h_m|s_m,s_{f_j})\sum_{k=0}^{j}\Delta_{s_m}\eta(s_{f_k}|s_m)$$

where

$$\sum_{k=0}^{j}\Delta_{s_m}\eta(s_{f_k}|s_m) = \sum_{k=0}^{j}\eta(s_{f_k}|s_{m+1}) - \sum_{k=0}^{j}\eta(s_{f_k}|s_m) \leq 0$$

if the cdf $H(s_f|s_m) \equiv \sum_{k=0}^{j}\eta(s_{f_k}|s_m)$ is decreasing in s_m—something we will verify below—so that higher s_m types are matched with higher s_f in the FOSD sense.

Thus, the first term in (O.3) is negative if $\Delta_{s_f}\Pi_m(h_m|s_m,s_f) \leq 0$ for all s_f (and given s_m); and the second term is negative if $\Delta_{s_m}\Pi_m(h_m|s_m,s_f) \leq 0$ for all s_m (and given s_f). To derive conditions under which this holds (i.e. under which $\Pi_m(h_m|s_m,s_f)$ is decreasing in s_m (given s_f) and in s_f (given s_m)) it suffices to show conditions under which the *marginal* pmf of male hours $\sum_{\tilde{h}_f}\pi_{(s_f,s_m)}(\tilde{h}_f,h_m)$ satisfies the monotone likelihood property (or, equivalently, is log-supermodular) in (h_m,s_m) (for fixed s_f), and in (h_m,s_f) (for fixed s_m), as these properties imply FOSD of the marginal cdf $\Pi_m(h_m|s_m,s_f)$ in male and female types, respectively.

The denominator of the joint pmf, $\pi_{(s_f,s_m)}(h_f,h_m)$, does not depend on the specific hours bundle (h_f,h_m), and so the denominator of the marginal pmf, $\sum_{\tilde{h}_f}\pi_{(s_f,s_m)}(\tilde{h}_f,h_m)$ does not depend on it either.

We can thus focus on the numerator of $\sum_{\tilde{h}_f} \pi_{(s_f,s_m)}(\tilde{h}_f, h_m)$ when establishing log-supermodularity. That is, we aim to show under which conditions:

$$\sum_{\tilde{h}_f} \exp(\overline{u}_{(s_f,s_m)}(\tilde{h}_f, h_m)/\sigma_\delta)$$

is log-supermodular in (h_m, s_m) (for fixed s_f), and in (h_m, s_f) (for fixed s_m). For this it suffices that the log-transformed summand is supermodular pairwise,[1] meaning in (h_m, s_m), (h_m, h_f), (h_f, s_m) (for fixed s_f) and in (h_m, s_f), (h_m, h_f), (h_f, s_f) (for fixed s_m), where

$$\log(\exp(\overline{u}_{(s_f,s_m)}(h_f, h_m)/\sigma_\delta)) = (w\,(\psi s_f h_f) + w(s_m h_m) + 2p^M(1 - h_m, 1 - h_f))/\sigma_\delta.$$

Supermodularity in (h_f, s_m) and in (h_m, s_f) is trivially satisfied (at equality). Supermodularity in (h_m, h_f) holds under the premise of strictly supermodular p. Finally, supermodularity in (h_m, s_m) and (h_f, s_f), holds if the wage function is supermodular in these pairs, which—as in the baseline model—is true if z is strictly supermodular and convex, as assumed.

A similar argument and analogous conditions establish that $\widehat{\Pi}_f(h_f|s_m)$ is decreasing in s_m. We have thus shown that under the premise of the proposition (and if marriage sorting is positive in a stochastic sense, to which we will turn next), labor hours are stochastically increasing in own and partner's type.

Part 3. Recall the probability that man s_m chooses woman s_f on the marriage market (see Section D.2):

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

We are interested in conditions under which:

$$\frac{\eta(s''_f, s''_m)}{\eta(s'_f, s''_m)} \geq \frac{\eta(s''_f, s'_m)}{\eta(s'_f, s'_m)}$$

for all $s''_f > s'_f$ and $s''_m > s'_m$. In words, we seek conditions under which probability $\eta(s_f, s_m)$ is log-supermodular, or equivalently, conditions that ensure the monotone likelihood ratio property of $\eta(s_f, s_m)$. The important implication will be that higher s_m-type men are matched to higher s_f-type women in the FOSD sense or, equivalently—using again the notation $H(s_f|s_m) \equiv \sum_{k=0}^{j} \eta(s_{f_k}|s_m)$ for the probability that a man with s_m marries a woman of type weakly below s_f—that $H(s_f|s_m)$ is decreasing in s_m.

Note that

$$\frac{\eta(s''_f, s''_m)}{\eta(s'_f, s''_m)} = \frac{\dfrac{\exp(\Phi(s''_m, s''_f, v(s''_f)))/\sigma_\beta)}{\sum_{\tilde{s}_f \in \{\mathcal{S} \cup \emptyset\}} \exp(\Phi(s''_m, \tilde{s}_f, v(\tilde{s}_f)))/\sigma_\beta)}}{\dfrac{\exp(\Phi(s''_m, s'_f, v(s'_f)))/\sigma_\beta)}{\sum_{\tilde{s}_f \in \{\mathcal{S} \cup \emptyset\}} \exp(\Phi(s''_m, \tilde{s}_f, v(\tilde{s}_f)))/\sigma_\beta)}} = \frac{\exp(\Phi(s''_m, s''_f, v(s''_f)))/\sigma_\beta)}{\exp(\Phi(s''_m, s'_f, v(s'_f)))/\sigma_\beta)}.$$

[1] This step uses the fact that log-supermodularity is preserved under discrete sums, as shown by de Clippel, Eliaz, and Rozen (2014), Proof of Theorem 4, Step 3.

Therefore, it suffices to show conditions that make $\exp(\Phi(s_m, s_f, v(s_f)))/\sigma_\beta$ log-supermodular, in (s_f, s_m) or, equivalently, conditions that make $\Phi(s_m, s_f, v(s_f))/\sigma_\beta$ supermodular in (s_f, s_m). Recall that:

$$\frac{\Phi(s_m, s_f, v(s_f))}{\sigma_\beta} = \frac{\sigma_\delta}{\sigma_\beta}\Big[\kappa + \log\Big(\sum_{h_f,h_m} \exp\Big\{\frac{w(s_m h_m) + w(\psi s_f h_f) + 2p^M(1-h_m, 1-h_f)}{\sigma_\delta}\Big\}\Big)\Big] - v(s_f)$$

And so $\Phi(s_m, s_f, v(s_f))/\sigma_\beta$ is supermodular in (s_m, s_f) if

$$\log\Big(\sum_{h_f,h_m} \exp\Big\{\frac{w(s_m h_m) + w(\psi s_f h_f) + 2p^M(1-h_m, 1-h_f)}{\sigma_\delta}\Big\}\Big)$$

is supermodular, or if

$$\sum_{h_f,h_m} \exp\Big\{\frac{w(s_m h_m) + w(\psi s_f h_f) + 2p^M(1-h_m, 1-h_f)}{\sigma_\delta}\Big\}$$

is log-supermodular in (s_f, s_m). This is the case if the summand

$$\exp\Big\{\frac{w(s_m h_m) + w(\psi s_f h_f) + 2p^M(1-h_m, 1-h_f)}{\sigma_\delta}\Big\}$$

is log-supermodular *pairwise*, or if $w(s_m h_m) + w(\psi s_f h_f) + 2p^M(1-h_m, 1-h_f)$ is supermodular pairwise in $(s_f, s_m), (s_f, h_f), (s_f, h_m), (s_m, h_f), (s_m, h_m), (h_f, h_m)$.[2] Thus, $\eta(s_f, s_m)$ is log-supermodular—and therefore higher s_f women match with higher s_m in the FOSD sense and thus on average—if the wage function is supermodular in (s_i, h_i) (which again is satisfied if z is supermodular and weakly convex) and if the home production function p is supermodular in (ℓ_m, ℓ_f). □

[2]This step again uses the fact that log-supermodularity is preserved under discrete sums, as shown by de Clippel, Eliaz, and Rozen (2014), Proof of Theorem 4, Step 3.

OC Data and Sample Construction

OC.1 Data Sources

In this section, we provide details on our three sources of data:

GSOEP. The main dataset used for the empirical analysis (Section 2) and the estimation (Section 4.4) is the German Socio Economic Panel (GSOEP), a household survey conducted by the German Institute of Economic Research (Deutsches Institut fuer Wirtschaftsforschung) starting in 1984. The core study of the GSOEP surveys about 25,000 individuals living in 15,000 households each year. All individuals aged 16 and older respond to the individual questionnaire. The head of household additionally answers a household questionnaire. This survey is longitudinal in nature, and collects rich information on demographics (such as marital status, education, fertility, family background, etc.), labor market variables (including hours worked, wages, and occupation), and detailed time-use information. Important for us, the GSOEP contains the same information for both the head of household and their partner (whether married or cohabiting).

Throughout our analysis, we focus on West Germany. Our baseline period is 2010-2016. We also consider an earlier period, 1990-1996, for over-time comparisons. We do not use data from before 1990 because key time-use variables are missing.

BIBB. Our main data source for measuring occupation types is the BIBB Employment Survey collected in 2012 by the German Federal Institute of Vocational Training (Bundesinstitut fuer Berufsbildung—BIBB), and the German Federal Institute for Occupational Safety and Health. The survey is representative of the German employed population. It contains data on task usage (self-reported by individuals on a discrete ordinal scale) for 1,235 occupations, defined at the 4-digit level (based on variable: `kldb92`). We merge this information at the occupational level with the GSOEP.

GTUS. For our empirical analysis on time use and its changes over time in Section 6, we complement the GSOEP with the German Time Use Survey (GTUS). The GTUS is administered by the German Federal Statistical Office. The survey assesses information on time use (work, leisure, home production, education, social engagements, training, etc.) and demographic characteristics of a representative sample of individuals and households, which allows us to link couples. It has three waves: 1991/92, 2001/02, 2012/13. In line with the time periods of our main sample from GSOEP, we focus on the first wave (1991/92) and the third wave (2012/13). The first wave surveys 6,400 households while the third one surveys 5,000 households.

The GTUS provides the most detailed and accurate time-use data for Germany. The data are collected using household and individual questionnaires and a time-use diary. The diary is filled out for several days: two days in 1991/92, and three days in 2012/13 (two weekdays and a weekend day). The time-use data is reported in short time intervals (5-minute intervals for 1991/92 and in 10-minute intervals for 2012/13). As our main goal is to illustrate properties of home production and their changes

over time, we focus on the time-use categories related to *home production*. We harmonize the available home production categories (Childcare, House Chores, Pets, Shopping, Household Organization, Meals, Textiles, Repairs, and Care, including commuting times for each category) across waves in order to make the analysis comparable over time. We aggregate the detailed time-use data to 'hours per day'.

We impose the same demographic restrictions as in our main GSOEP sample (see Section OC.2.1).

OC.2 Sample for Empirical Facts

We now describe the sample restrictions and the variables used for our empirical facts in Section 2.

OC.2.1 Sample restrictions

For our *Main sample*, we pool observations from the period 2010-2016, from the original GSOEP samples and their refreshments.[3] For most of the analysis, we restrict our attention to West Germany.[4]

We impose the following demographic restrictions: We keep all individuals in private households, either singles or heterosexual couples (married or cohabiting). We restrict our analysis to individuals who either never married or are in their first marriage.[5] We keep individuals who are in their prime working age, 22-55 years old.

Regarding the labor market, we exclude from our sample those individuals who are self-employed or still in school, those working in odd occupations (identified with the occupational code `kldb92`$\geq$ 9711), and those who are employed but with missing occupational code.

We impose these restrictions at the individual level. This implies that when we analyze individual outcomes, one partner of a given couple could be in the sample, while the other partner is not. For the analysis of couple outcomes, we keep couples in which both partners fulfill our sample restrictions.

OC.2.2 Variable description

We now describe the variables we use in the empirical analysis.

1. **Education:** We classify individuals into three education groups: a) 'Low Education' includes individuals with either only a high school degree (13 years of schooling) or those with middle school degrees (9 or 10 years) plus some basic vocational training ($<$ 11 years of schooling); b) 'Medium Education' includes those with either high school or middle school degrees and vocational training, with $>=$ 11 years of schooling; c) 'High Education' includes those with a college degree or more. Education levels are defined based on the ISCED-97 classification. Alternatively, we use years of education as our schooling measure, and we left-truncate this variable at 10 years of education.

[3] We exclude from our analysis the migrants and refugees samples, the oversampling of low income individuals and single parents, and the oversampling of high income earners.

[4] We drop Berlin from the sample since it cannot be unambiguously assigned to East or West Germany. This is standard in the literature, see for example Heise and Porzio (2019).

[5] Since cohabitants are defined as never married, we cannot rule out that we capture a cohabiting relationship that is not the first one.

2. **Marriage Market Sorting:** For our graphical analysis at the individual level, we define marriage market sorting bins by the difference between the years of education of an individual and the years of education of their partner. For the graphical analysis at the couple level, we define marriage market sorting bins as the difference between the years of education of the male partner and the years of education of the female partner. Beyond the graphical analysis, we measure marriage market sorting as the correlation between the education levels of partners.
3. **Labor Market Sorting:** We define labor market sorting as the correlation between the individual's years of education and their matched job characteristic—the task complexity of the individual's occupation (defined in Appendix OD.2).
4. **Hours:**
   (a) **Labor Market Hours**: We define labor market hours as the number of self-reported hours that an individual works in a given week (including overtime). We winsorize hours to 10 and 60 hours, at the bottom and the top, respectively.
   (b) **Home Production Hours:** We measure home production hours as the weekly time an individual allocates to the following activities: childcare, housework (which includes household chores such as cooking, cleaning, etc.), running errands, repairs of the house or car, and garden work. Since home hours are measured on a typical weekday, we multiply them by five, for consistency with market hours. We impute missing data on home hours using information on labor hours, assuming a total weekly time budget of 70h for work at home and in the labor market.
   (c) **Leisure Hours:** We measure leisure hours as the weekly hours allocated to hobbies and other leisure activities (measured as daily leisure hours $\times$ 5).

OC.3 Estimation Sample

In this section, we describe the sample restrictions and variables in our estimation sample.

OC.3.1 Sample Restrictions

In order to construct our *Baseline Estimation Sample*, we use data from West Germany for the period 2010-2016 of the original GSOEP and its refreshments (as discussed above). For our *Past Estimation Sample*, we use data from West Germany for the period 1990-1996.[6] We drop those individuals that appear in both periods. We then apply the following restrictions, similar to our *Main Sample* above:

1. **Age Restrictions**: We restrict our attention to individuals between 22 and 55 years old. We keep couples in which both partners are within this age range.[7]
2. **Marital Status Restrictions**: We focus on individuals who are either single or in heterosexual couples (married or cohabiting). We restrict our analysis to their first marital spell, as defined

[6] We drop individuals observed both in West and East Germany while in their typical occupation, see below.
[7] We drop married individuals whose spouses' information is missing during the entire sample period.

in OC.2.1. We drop observations from periods after the first marriage ended, or for which the end date of the first marriage cannot be identified. We drop individuals for whom we can identify more than one spouse/partner during the sample period.

3. **Labor Market Restrictions:**
   (a) We exclude from our sample observations corresponding to individuals working in odd occupations (`kldb92`$\geq$ 9711) or that are employed but have missing occupation codes.
   (b) We drop employed individuals with missing data on hourly wages.
   (c) We drop the self-employed and individuals who are still in school, as defined by their typical occupation (see Online Appendix OC.3.2).
4. **Additional Restrictions:**
   (a) We exclude observations from individuals to whom we cannot assign a human capital type. The estimation of human capital types is discussed in Appendix OD.1.
   (b) We drop individuals in couples for which information on the spouse/partner is always missing.

OC.3.2 Definition of Typical Occupation, Hours, Wages and Marital Status

In this section, we explain our methodology to create measures of individual-level labor market and demographic variables (which we refer to as 'typical outcomes'). This allows us to bridge the dynamic features of the data with the static nature of our model.

Typical Occupation: We use the following rules to assign a typical occupation to an individual.

1. If they appear in the sample only once, or if they appear more than once but always report the same occupation, we assign to them that unique occupation.
2. If they appear in the sample more than once, and in at least one of those years they are in a *'non-labor market state'* (i.e., 'self-employed', 'studying' or 'not-employed'), we proceed as follows:
   (a) If they were either 'not-employed' or 'self-employed' or 'studying' for strictly more than half of the time they appear in the sample, we consider that state as their typical occupation.
   (b) If they were in one of these states (e.g., 'self-employed') exactly half of the time they appear in the sample, and spent the other half in the other two states (e.g., 'studying' and 'not-employed'), we assign them to the state in which they spent half the time ('self-employed').
   (c) If (a) and (b) do not hold, but they spent more than 75% of their time in the 'non-labor market states' combined, then we assign them to the state with the longest duration.
3. If we observe an individual multiple times, but only in a single occupation (and they hold this occupation for more than 25% of the time), we assign them to that unique occupation.
4. If we observe an individual in more than one occupation during the sample period, we construct the difference in percentiles of task complexity between their highest and lowest-ranked occupation (where occupations are ranked as described in Appendix OD.2). We then proceed as follows:
   (a) When the difference in percentiles is ≥ 0.1, we assign them to their highest ranked occupation.
   (b) If the difference is < 0.1, we assign them to the occupation with the longest tenure. If there

is a tie, we assign them to the highest ranked occupation among those with equal tenure.

5. After applying these rules, we drop from the sample individuals whose typical occupation is 'self-employment' or 'studying'.

Typical Labor Market Hours: We define typical labor market hours as the average self-reported work hours (including overtime) over the years an individual was in their typical occupation, defined above. We winsorize labor hours to 10 and 60 hours, at the bottom and the top, respectively. For those individuals whose typical occupation is 'not employed', typical labor hours are set to zero.

Typical Hourly Wage: For each individual, we define their typical wage as the average real hourly wage over the years they worked in their typical occupation.[8]

Typical Home Production Hours: We construct typical home production hours (defined in Online Appendix OC.2) as the average home production hours over the years an individual was in their typical occupation. We impute missing data for home hours as discussed in Online Appendix OC.2.

Typical Marital Status: We define the typical marital status based on the following rules:

1. If the individual had only one marital status during the sample period, we consider that marital status as the typical one.
2. If the individual switched from being single to being married during the sample period, we assign the marital status observed when employed in their typical occupation. If they were observed as both single and married while in their typical occupation, we assign them to marriage.
3. We exclude from the sample individuals that report more than one spouse during the sample period. We also exclude their partners.

Typical Child: In Section 6/Appendix G.4, we re-estimate our model on two different subsamples: those with and those without children. To assign individuals to either subsample, we define their 'typical child' status based on the presence of children under 18 years old in the household, during at least one of the years in which they are in their typical occupation. If there are discrepancies between husband and wife (4% of cases), we classify households based on the 'typical child' status of the male partner. For singles, we randomly distribute them to either subsample so that the share of singles in each subsample equals the share of singles in the whole sample.

Following these rules, our *Baseline Estimation Sample* (West Germany, 2010-2016) has 3,857 individuals living in 2,326 households. Of these households, 1,531 are couples and 795 are singles (418 are single women and 377 are single men).

Our *Past Estimation Sample* (West Germany, 1990-1996) consists of 2,336 individuals in 1,294 households, of which 1,042 are couples and 252 are singles (117 single women and 135 single men).

[8] We construct hourly wages based on inflation adjusted monthly earnings, divided by monthly hours (constructed as weekly hours times 4.3). Hourly wages are trimmed at the bottom and top 1% percentile. The data for inflation adjustment comes from the OECD: https://data.oecd.org/price/inflation-cpi.htm

OD Estimation of Worker and Job Types

OD.1 Estimation of Worker Types

OD.1.1 Sample Selection, Methodology and Key Variables

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 corresponding to self-employment or work in odd occupations (`kldb92`$\geq$ 9711). We also exclude observations with missing information on education or with missing (not zero) labor force experience. Our panel consists of around 212,000 person-year observations.

Methodology. 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

$$\ln w_{it} = \nu_i + \sum_{ed \in \{med, high\}} \alpha^{ed} x_{it}^{ed} + \beta_1 h_{it} + \beta_2 h_{it}^2 + \beta_z' Z_{it} + \delta_s + \delta_t + \epsilon_{it}, \tag{O.4}$$

where x_{it}^{ed} are indicator variables for the education group of an individual (capturing x in our model). Coefficient α^{ed} gives the return to education $ed \in \{med, high\}$ in terms of log wage units, where $0 < \alpha^{med} < \alpha^{high}$ would indicate positive returns to education. In turn, ν_i is an individual fixed effect that captures unobserved time-invariant ability, with model counterpart ν. Further, h_{it} denotes weekly labor hours (capturing h in our model). Finally, Z_{it} are time-varying controls for the individual, δ_s and δ_t are state and time fixed effects, and ϵ_{it} is a mean-zero error term.[9] We implement (O.4) as a panel wage regression. Below we explain how we deal with the following three challenges in the estimation: selection into labor force participation, endogeneity of hours worked, and agents who hold jobs for less than two years in our sample.

Key Variables. For *weekly hours*, we use reported actual hours, see Online Appendix OC.2.2. For *labor force experience*, we use the reported labor force experience and we impute it by potential experience if this information is missing. The three *education* groups (low, medium, high) are defined in Online Appendix OC.2.2. Our wage variable is *log hourly wages*, inflation-adjusted in terms of 2016 Euros. For the definition of 'demographic cells' in the selection stage, we additionally use the presence of *children below 3 years old* in the household, *age bins* (≤ 25, > 25 and ≤ 40, > 40 and ≤ 50, > 50) and *state* of residence.

[9] We do not include occupation fixed effects since in our model, conditional on $\tilde{s}$ (which we control for by including regressors (x, ν, h)), the wage does not depend on occupation y 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.

OD.1.2 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 (O.4). Since the variation in participation in our sample is mainly driven by women, we use the 'progressiveness' in an individual's narrowly defined demographic cell, proxied by the share of females working in that cell. Our cells are defined by a combination of state of residency, year, age bin, and an indicator of whether a child under the age of 3 is in the household.[10] This strategy builds on the literature on grouping instruments (Angrist, 1991 and Blundell, Duncan, and Meghir, 1998) which proposes to assign individuals to groups and use group-specific averages to instrument for endogenous variables at the individual level.[11] When defining this variable for a particular individual, we employ the 'leave-one-out' method to avoid counting the individual's labor force participation when computing this statistic. We further drop cells with less than five observations, yielding around 2,500 cells. 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.

While in our empirical selection equation, labor force participation is driven by the progressiveness of the demographic cell, in the model this selection is determined by labor supply shocks. We argue, however, that these sources of selection are related. Low realizations of the labor supply shocks that lead to non-participation in the model partly reflect obstacles to the labor force participation of women. These hurdles are likely specific to locations and demographic groups, and can stem from institutional factors (e.g., the low regional availability of childcare slots); or from supply side forces (e.g., low availability of immigrant workers, as in Cortés and Tessada, 2011); or from conservative norms and culture prevailing in a location and therefore among the peers whose opinions affect own labor force participation choices (see Fernández and Fogli, 2009, Maurin and Moschion, 2009, Fogli and Veldkamp, 2011, Mota, Patacchini, and Rosenthal, 2016, Fortin, 2015, Olivetti, Patacchini, and Zenou, 2020). To the extent that these barriers affect labor supply choices of most women in a certain location/demographic group, our IV is a good proxy for the labor supply shocks in the model. Specifically, we think of the individual's labor supply shock in our model as being positively correlated with our progressiveness IV in the data, bridging selection in model and data.

[10] German children start kindergarden when they are three years old. Before this age, children are predominantly at home (in 2013, only 29% of children aged 0-2 were in daycare (OECD, 2016)), so age three is an important threshold that affects maternal labor force participation.

[11] The idea is that group-specific means capture the aggregate/institutional *common* environment a group faces, which explains members' behavior but is uncorrelated to unobserved idiosyncratic errors. In light of this literature, our cell definition is closest to that in Kalíšková (2020) who studies the impact of tax benefits on female labor supply. To instrument for individual participation tax benefits she employs the average tax benefits when out of work, by education, presence of children by age, marital status, country and year. Relatedly, Burns and Ziliak (2016) identify the impact of marginal tax rates on taxable income, using cohort-education-state average tax rates to instrument for person-specific tax rates.

We run the following probit selection regression:

$$emp_{it} = \alpha\, share_{jit} + \sum_{ed \in \{med, high\}} \alpha^{ed} x_{it}^{ed} + \beta_z' Z_{it} + \delta_s + \delta_t + \epsilon_{it} \tag{O.5}$$

where the dependent variable indicates whether individual i is employed at time t. Variable $share_{jit}$ 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 above), x_{it}^{ed} is an education indicator ('low education' is the reference group), Z_{it} is a vector of individual controls (linear and quadratic labor force experience in years, household size). Finally, δ_s and δ_t are state and year fixed effects and ϵ_{it} is a mean-zero error term.

The results are in Table O.1, where we label the variable $share_{jit}$ by *Share of Working Women in Cell*. We find a strong positive effect of the share of women working in a demographic cell on labor force participation of an individual in that cell.

Table O.1: Selection Regression

	Employed
Share of Working Women in *Cell*	1.355***
	(0.029)
Experience	0.092***
	(0.002)
Experience^2	-0.002***
	(0.000)
Medium Educ	0.453***
	(0.009)
High Educ	0.851***
	(0.012)
HH Size	-0.046***
	(0.0054)
Constant	-0.918***
	(0.0358)
State and Year FE	Yes
Observations	212,894
Pseudo R^2	0.144

Notes: Standard errors clustered at cell level in parentheses. ***p<0.01, **p<0.05, *p<0.1. Results are based on regression (O.5). A *Cell* is defined by state, age, year and presence of children under 3 in the household.

OD.1.3 Panel Wage Regression

One potential issue for our specification of the wage regression is that there may be unobserved 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 (see Section 2 for a discussion) could still be problematic. To address this concern, we instrument own labor hours by partner's labor hours in regression (O.4). More specifically, we instrument own labor hours and labor hours squared by (i) labor hours of the partner, (ii) labor hours of the partner squared and (iii) whether a partner is present (where we set the partner's hours to zero in two cases, either if the partner is present but not working or if no partner is present).

Through the lens of our model this instrument is valid: First, there is a systematic relationship between the hours worked of an individual and the hours worked by their partner. Second, the partner's labor hours impact own wage *only* through own labor hours.

We drop observations for individuals whose partner reports to be employed but has zero reported labor hours, or whose partner has missing employment and hours information. Since we include individual fixed effects we also drop observations for individuals who are in our panel for only a single year. We further restrict the sample to those who are employed and have non-missing hourly wages. This leads to a sample of 133,214 person-year observations. Based on our model wage function, we choose the following regression specification:

$$\ln w_{it} = \nu_i + \sum_{ed \in \{med, high\}} \alpha^{ed} x_{it}^{ed} + \beta_0 IMR_{it} + \beta_1 h_{it} + \beta_2 h_{it}^2 + \beta_z' Z_{it} + \delta_s + \delta_t + \epsilon_{it} \tag{O.6}$$

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 (O.5), h_{it} are weekly hours worked, and x_{it}^{ed}, Z_{it}, δ_s and δ_t are as in (O.5). 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 (O.6).

Table O.2 contains the results. Columns (1) and (2) report the first stage regressions for two variables to be instrumented: weekly hours and weekly hours squared. Column (3) contains the results of the second stage regression. Based on the F-statistics, 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. 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 positive effect on hourly 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 around 4%, computed from the estimated coefficients in Column (3) as $(40 \times 0.119 - 40^2 \times 0.00164) - (30 \times 0.119 - 30^2 \times 0.00164) = 0.042$.

OD.1.4 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 covariates that are most correlated with the individual fixed effects (such as *education*, *gender* and *full time labor force participation*). After imposing the sample restrictions described in Online Appendix OC.3.1, our final estimation sample (for period, 2010-2016) has 3,857 unique individuals, 24% of them have imputed ν_i.

Table O.2: Wage Regression

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

Notes: *** p<0.01, ** p<0.05, * p<0.1. Standard errors clustered at the *Cell* level are in parentheses, where *Cell* is defined by state, year, age, and presence of children under the age of 3 in the household. The *Inverse Mills Ratio* is defined for each individual i in year t and based on the selection probit regression (O.5). Column (3) reports the IV version of regression (O.6), where we instrument *Weekly Hours Worked* and *Weekly Hours Worked*^{2} by the partner's weekly hours worked (linear and squared) and whether the partner is present. The results of the first stage are in columns (1) and (2).

OD.1.5 Distribution of Workers' Human Capital

We then divide individuals into our three education bins and assess within each bin whether an individual has a low (below median) or high (above median) fixed effect, so there are two subgroups $j \in \{1, 2\}$ in each education bin. We compute the subgroup fixed effect, $\overline{\nu}_j^{ed}$, as the mean of the individual fixed effects in subgroup j of education bin ed. Finally, we compute the human capital type for each individual i as $s_i = \alpha^{ed} x_i^{ed} + \overline{\nu}_{ji}^{ed}$ where $\overline{\nu}_{ji}^{ed}$ is the mean fixed effect of individual i's group j. We obtain six s-types. The resulting human capital distribution (s-types by education group) is displayed in Table O.3.

Table O.3: Worker Distribution of s-Types by Education

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

Notes: The education groups (*Low, Medium, High*) are defined in Online Appendix OC.2.2. The table displays six s-types, ordered by their value of $s = \alpha^{ed} x^{ed} + \overline{\nu}_j^{ed}$ for $j \in \{1, 2\}$ and $ed \in \{low, med, high\}$. Parameters α^{ed} are estimated in regression (O.6) and $\overline{\nu}_j^{ed}$ is the mean of individual fixed effects in education group ed and subgroup j, also based on (O.6).

OD.2 Estimation of Job Types

Sample Selection. Our main data source for measuring occupation types is the BIBB, see Online Appendix OC.1 for details. In order to reduce the problem of noisy reporting, we drop occupations for which task information is based on less than 5 individuals. We are left with task data for 613 occupations. Our structural estimation exercise will be based on the subset (608 occupations) that can be merged to the occupations held by individuals in our GSOEP sample, using the four-digit occupational codes `kldb92` from the German Classification of Occupations 1992.

Task Data. The BIBB contains data on how intensely different tasks are performed at each of our 4-digit occupations. The set of tasks is 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. See Online Appendix OC.1 for details.

Model Selection Stage. We merge the task data from the BIBB into occupations held by individuals in the GSOEP. As with the estimation of 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 except: Improve Process, Consequence of Mistakes and Difficult Decisions).

Principal Component Analysis (PCA). 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 (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 (43%). Based on the loadings on the underlying task descriptors, we interpret this component as *task complexity*: all task variables have positive loadings except *detailed work* and *same cycle*, which arguably are the only tasks in the dataset that indicate routine work. Figure O.3 shows a scree-plot with eigenvalues of the different principal components (left) and the loadings of the first principal component (right).

Figure O.3: Principal Component Analysis

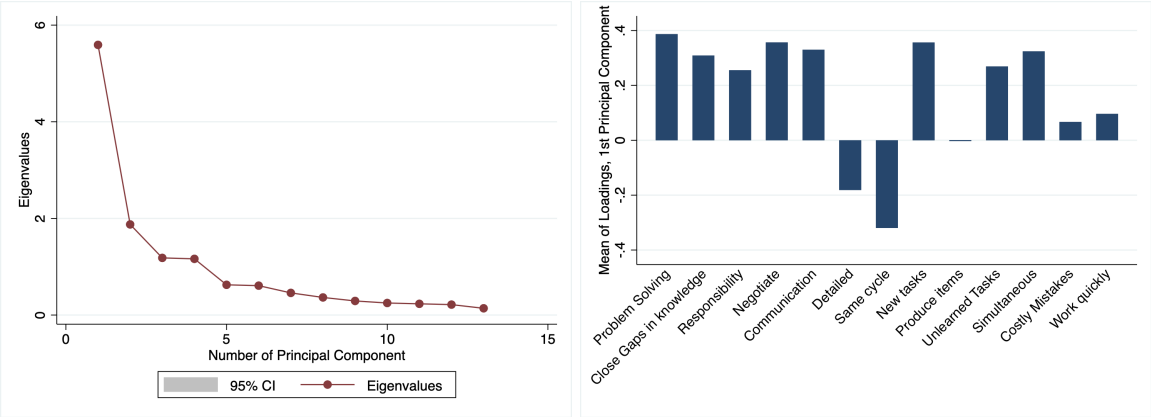

Distribution of Jobs' Task Complexity. We then compute the mean of this first principal component by occupation and denote it by $\hat{y}$. Once matched to our main sample (see Online Appendix OC.3), we define our final measure of occupational or job type as the *rank* of the occupation in the task complexity distribution, i.e. $y = \widehat{G}(\hat{y})$, where $\widehat{G}$ is the cdf of $\hat{y}$. So $y \sim G$ where $G = U[0, 1]$. We use this transformation because the occupational task data has an ordinal interpretation. Examples of occupations in the top 5% of the G distribution include engineers and programmers. Examples of occupations in the bottom 5% include janitors and cleaners.

OE Estimation: Construction of Moments

Our main estimation targets 16 moments defined in Table O.4. In this section, we provide details on how we construct these moments in the data and in the model.

For moment M1, we compute the female to male labor force participation ratio, including both married and single individuals. Both in the data and in the model, we define labor force participation as a dichotomous variable that takes value 1 when an individual works positive hours, and zero otherwise.

For moment M2, we compute the ratio of female full-time workers to male full-time workers, including both married and single individuals. Both in the data and in the model, we define 'full-time' as working more than 44% of the available time. This is equivalent to working more than 37.5 hours per week in the data, and to working more hours than described by the fifth entry in the hours grid in the model.[12]

For M3 and M4, we compute the married to single ratio in labor force participation, separately for women and men.

We compute M5 as the correlation between the time that female and male partners spend in home production, where home production hours are constructed as the share of total available time spent in home production.[13] We restrict our attention to individuals in couples.

To construct wage moments M6-M9, we use data on hourly wages for all employed individuals, whether they are single or in a couple. See Table O.4 for more details.

Moment M10 is constructed as the correlation of partners' s-types (see Appendix OD.1 for the details on the human capital estimation).

Moments M11 and M12 measure the gender wage gap by (s, h)-types. We use two (s, h)-type combinations: all individuals of either s-type 3 or s-type 4 (see columns 3 and 4 of Table O.3) that work full-time in the labor market.

For M13-M14, we compute the labor force participation rate of women in couples where both partners have a similar human capital type. For M13, we pool couples in which both partners are either of s-type 3 or 4 (corresponding to columns 3 and 4 in Table O.3). For M14, we pool couples in which both partners are either of s-type 5 or 6 (corresponding to columns 5 and 6 in Table O.3).

Finally, for M15 and M16, we compute the labor force participation rate among single women. M15 considers single women of s-typez 3 or 4. M16 pools single women of s-types 5 or 6.

[12] In the data, we determine the total available time by the 95th percentile of the distribution of total time spent working (home production plus labor market work). This is 85 weekly hours for the period 2010-2016.

[13] We define the total available time as in footnote 12 to deal with the fact that different individuals report different total hours allocated to home production and market work. When individuals report more than 85 hours of home production per week, we say that their share of time allocated to home production equals 1.

Table O.4: Moments

Moment Description	**Definition**
Labor Force Participation Female to Male Ratio (M1)	$\frac{Pr(h_f>0)}{Pr(h_m>0)}$
full-time Work Female to Male Ratio (M2)	$\frac{Pr(h_f=\hat{h})}{Pr(h_m=\hat{h})}, \hat{h} \geq 37.5$
Labor Force Participation Married to Single Ratio, by Gender (M3-M4)	$\frac{Pr(h_i>0\|Married)}{Pr(h_i>0\|Single)}, i \in \{f, m\}$
Correlation of Spouses' Home Production Hours (M5)	$corr(1-h_f, 1-h_m)$
Mean Hourly Wage (M6)	$\mathbb{E}[w]$
Variance of Hourly Wage (M7)	$Var[w]$
Overall (90-10) Wage Inequality (M8)	$\frac{w^{90}}{w^{10}}$
Upper Tail (90-50) Wage Inequality (M9)	$\frac{w^{90}}{w^{50}}$
Correlation between Spouses' Human Capital Types (M10)	$corr(s_m, s_f)$
Gender Wage Gap by Effective Type (M11-M12)	$\frac{\mathbb{E}[w(h_i s_i)\|i=m,h_i=\hat{h},s_i=\hat{s}]-\mathbb{E}[w(s_i h_i)\|i=f,h_i=\hat{h},s_i=\hat{s}]}{\mathbb{E}[w(h_i s_i)\|i=m,h_i=\hat{h},s_i=\hat{s}]}$
Female Labor Force Participation by Couple Type (M13-M14)	$Pr(h_f > 0\|s_f = s_m = \hat{s})$
Female Labor Force Participation of Single Women by Type (M15-M16)	$Pr(h_f > 0\|Single, s_f = \hat{s})$

OF Comparative Statics with Respect to Key Drivers of Inequality

To highlight the main forces behind inequality, we conduct comparative statics exercises in our estimated model with respect to the key drivers of inequality (ρ, θ, ψ).

The Effect of ρ. We are interested in the effects on inequality when ρ becomes even more negative (i.e., if home hours become even more complementary).

Figure O.4, first row, plots the effect of ρ on different inequality measures: 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, as the complementarity in home production 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) and thus productivity, which puts downward pressure on the wage gap. Moreover, because women 'improve' a sorting-relevant attribute (labor hours) relative to men, the gender gap in labor market sorting declines (panel f), reducing the gender wage gap even further.

Figure O.4(c) then 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 stemming from stronger marriage sorting.

The Effect of θ. We are interested in the effect of women's relative productivity at home on inequality. Eliminating the gender gap in home productivity (reducing θ from the estimate $\theta = 0.82$ to $\theta = 0.5$) would cut the gender wage gap by almost half, see Figure O.5(a): Making home productivity more equal increases the incentive for positive marriage sorting (panel d) and pushes against 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 O.5(c). Here, this is driven by a decrease in within-household inequality, which dominates the increase in between-household inequality driven by a rise in marriage sorting (Figure O.5(b)).

The Effect of ψ. Last, Figure O.6(a) shows that eliminating the female labor market wedge ψ (i.e., increasing ψ from our estimate $\psi = 0.85$ 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. Finally, Figures O.6(b) and (c) show that 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.

In sum, we derive several insights: First, eliminating asymmetries in productivity across gender (whether at home through $\theta \to 0.5$ or at work through $\psi \to 1$) reduces the gender wage gap. But this is not the only way to mitigate 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 gender gaps in labor hours and labor market sorting 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 O.4: 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.

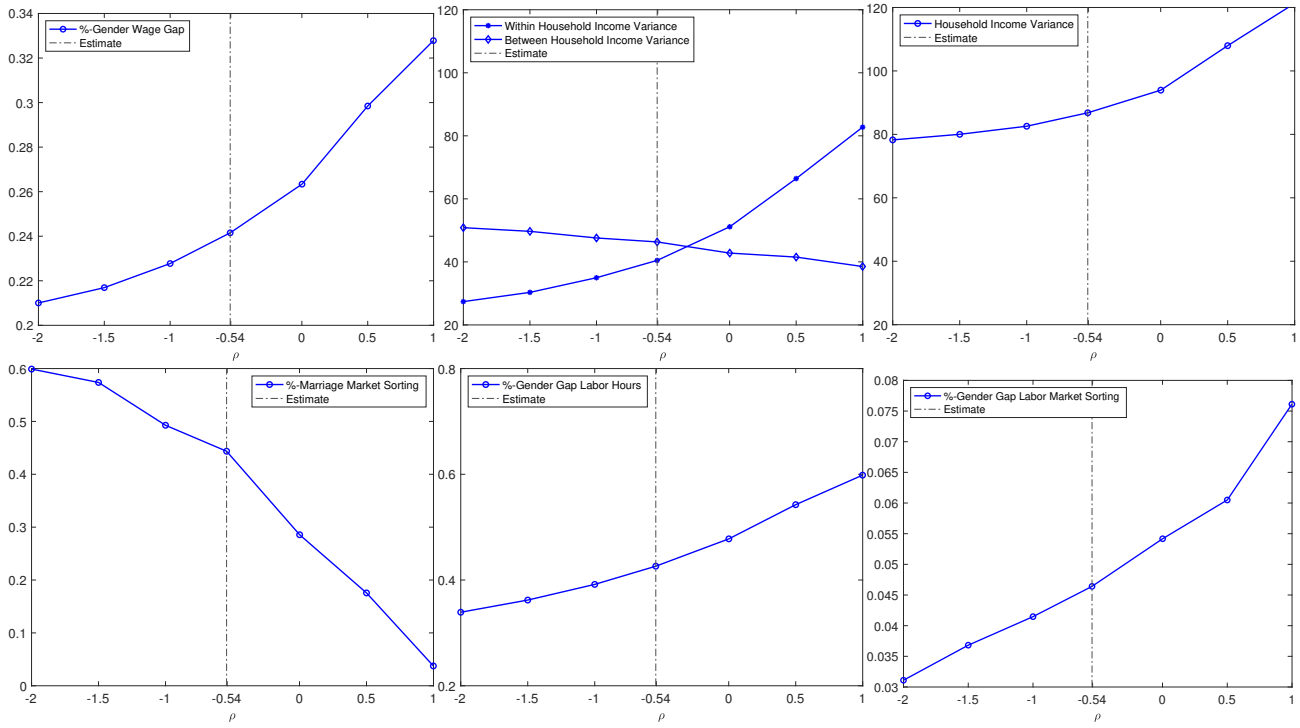

Figure O.5: a. Gender Wage Gap, b. Household Income Variance Decomposition, c. Household Income Variance, d. Marriage Market Sorting, e. Gender Gap Labor Hours, f. Gender Gap Labor Sorting.

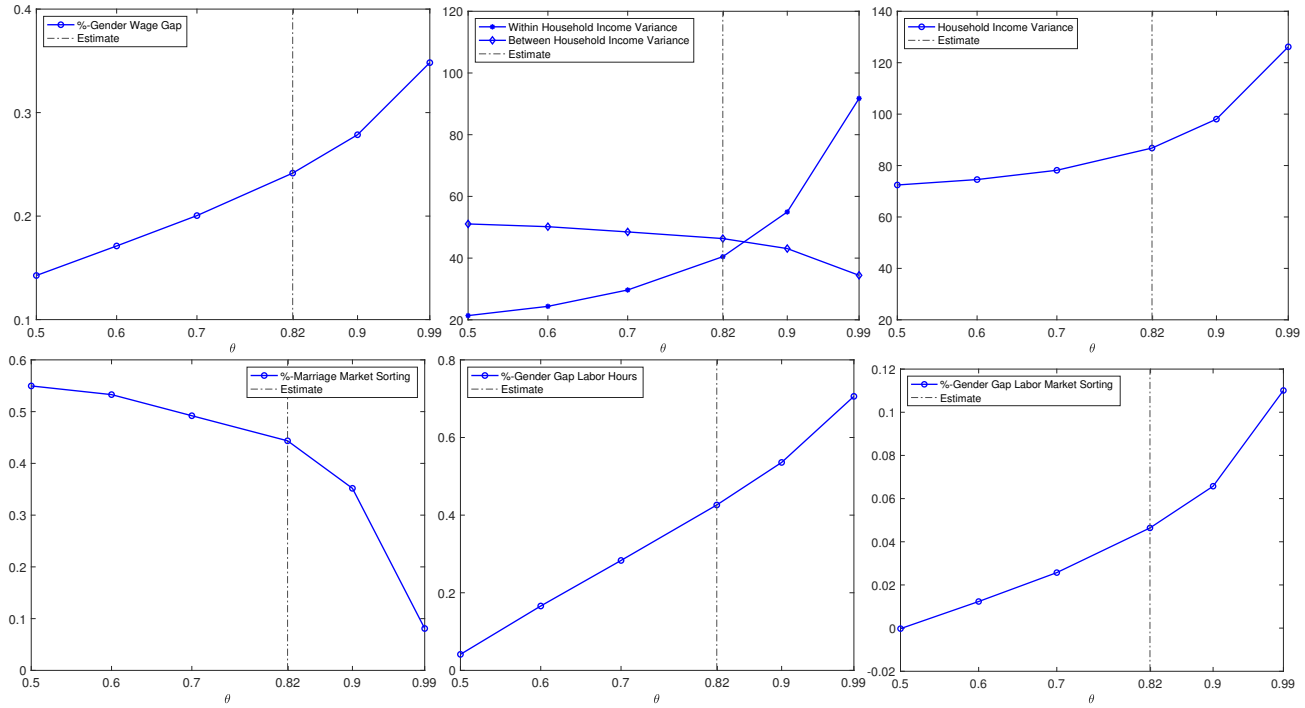

Figure O.6: a. Gender Wage Gap, b. Household Income Variance Decomposition, c. Household Income Variance, d. Marriage Market Sorting, e. Gender Gap Labor Hours, f. Gender Gap Labor Sorting.

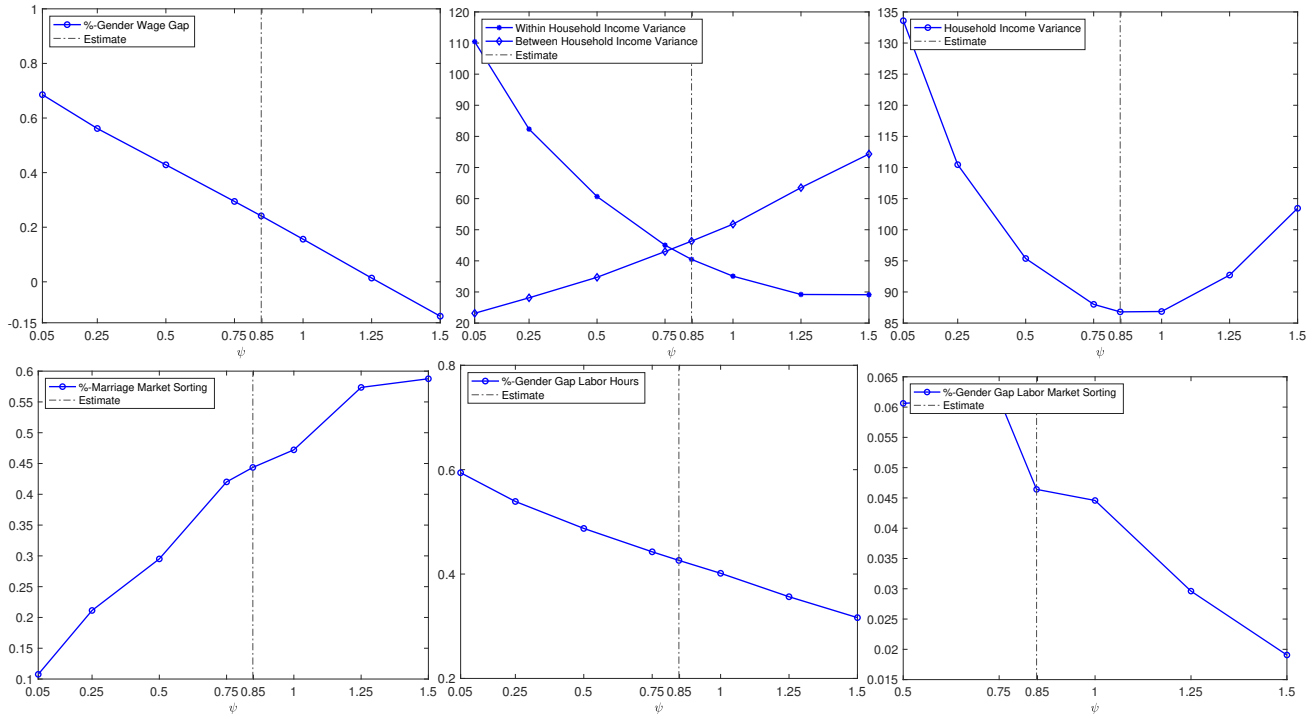

OG Home Production in GTUS vs. Home Production in GSOEP

Even though home production tasks in the GSOEP are less disaggregated, we find similar patterns as in the GTUS. The magnitude of the over-time change in aggregate home production correlation is similar across datasets, from 0.09 to 0.24 between 1991/92 and 2012/13 in GTUS and from 0.19 to 0.32 between 1990-1996 and 2010-2016 in GSOEP. Moreover, also in the GSOEP, partner's correlation in childcare hours is larger than in other home production activities (0.53 versus 0.07). This also holds when controlling for confounding factors, see Table O.5 below.

Table O.5: Complementarity in Home Production Hours: Childcare vs. Housework

	(1)	(2)
	Childcare Male Hours	Housework Male Hours
Childcare Female Hours	0.207***	
	(0.024)	
Housework Female Hours		0.117***
		(0.021)
Demographic Controls	Yes	Yes
State and Year FE	Yes	Yes
Period	2010-2016	2010-2016
Observations	4,007	3,874
R-squared	0.215	0.080

Notes: Standard errors clustered at the state level in parentheses. *** p<0.01, ** p<0.05, * p<0.1. We impose the same sample restrictions and controls as in Table A.6. We further restrict our attention to the sample of households with children under 18 years old.

References

J. D. Angrist. Grouped-data estimation and testing in simple labor-supply models. *Journal of Econometrics*, 47(2):243–266, 1991.

R. Blundell, A. Duncan, and C. Meghir. Estimating labor supply responses using tax reforms. *Econometrica*, 66(4):827–861, 1998.

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

S. K. Burns and J. P. Ziliak. Identifying the elasticity of taxable income. *The Economic Journal*, 127 (600):297–329, 2016.

Y. Chang and F. Schorfheide. Labor-supply shifts and economic fluctuations. *Journal of Monetary Economics*, 50(8):1751–1768, 2003. ISSN 0304-3932. doi: https://doi.org/10.1016/j.jmoneco.2003.02.001. URL https://www.sciencedirect.com/science/article/pii/S0304393203001132.

P.-A. Chiappori and E. Gugl. Transferable utility and demand functions. *Theoretical Economics*, 15(4), 2020.

P. Cortés and J. Tessada. Low-skilled immigration and the labor supply of highly skilled women. *American Economic Journal: Applied Economics*, 3(3):88–123, July 2011.

G. de Clippel, K. Eliaz, and K. Rozen. Competing for consumer inattention. *Journal of Political Economy*, 122(6):1203–1234, 2014. ISSN 00223808, 1537534X. URL http://www.jstor.org/stable/10.1086/677253.

R. Fernández and A. Fogli. Culture: An empirical investigation of beliefs, work, and fertility. *American Economic Journal: Macroeconomics*, 1(1):146–77, 2009.

A. Fogli and L. Veldkamp. Nature or nurture? Learning and the geography of female labor force participation. *Econometrica*, 79(4):1103–1138, 2011.

N. M. Fortin. Gender role attitudes and women's labor market participation: Opting-out, aids, and the persistent appeal of housewifery. *Annals of Economics and Statistics*, (117/118):379–401, 2015.

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

S. Heise and T. Porzio. Spatial wage gaps in frictional labor markets. 2019.

K. Kalíšková. Tax and transfer policies and the female labor supply in the EU. *Empirical Economics*, 58(2):749–775, 2020.

E. Maurin and J. Moschion. The social multiplier and labor market participation of mothers. *American Economic Journal: Applied Economics*, 1(1):251–72, January 2009.

N. Mota, E. Patacchini, and S. Rosenthal. Neighborhood effects, peer classification, and the decision of women to work. Working Paper, 2016.

OECD. Starting strong IV. Early childhood education and care data country note. Germany. Technical report, OECD Publishing, Paris, 2016.

C. Olivetti, E. Patacchini, and Y. Zenou. Mothers, peers, and gender-role identity. *Journal of the European Economic Association*, 18(1):266–301, 2020.