The Labor Hub

Job Market Paper Series No. 17

Advisor: Pengpeng Xiao

FERTILITY, MARRIAGE MARKET EQUILIBRIUM, AND EDUCATION CHOICE

Frederik Almar and Yusuke Ishihata

Fertility, Marriage Market Equilibrium, and Education Choice *

Frederik Almar †

Yusuke Ishihata‡

Job market paper

October 26, 2026

[Click here for the latest version]

Abstract

Fertility in Denmark has fallen along with partnership rates, even though educational attainment has risen and is positively correlated with fertility. To understand the mechanisms behind these patterns, we develop and estimate an equilibrium model of education, marriage, labor supply, and fertility with imperfectly transferable utility and unobserved heterogeneity in fertility preferences. We identify women's and men's preferences for children separately from couples' joint choices using variation in local marriage market conditions that shifts bargaining power. The estimated model attributes the decline mainly to falling preferences for children among women, while men's preferences change little. Economic factors such as wages and the difficulty of combining work and children do not explain the decline. Policy counterfactuals show that holding education, marriage, and bargaining power fixed can substantially understate the effects of policies on fertility. Free childcare raises fertility by 3.4 percent, and partnership formation accounts for 41 percent of that gain. Closing the gender wage gap reduces fertility in part through a shift in bargaining power toward wives, who want fewer children than their husbands.

Keywords: Fertility, Marriage Market, Education Choice, Bargaining Power, Policy
JEL Classification: D13, I24, J12, J13, J16, J18

*Ishihata is grateful to his advisors Peter Arcidiacono, Pengpeng Xiao, Arnaud Maurel, and Laura Pilossoph for their continued guidance and support. We thank Jérôme Adda, Jason Baron, Pat Bayer, Edoardo Ciscato, Michael Dinerstein, Kenta Fukuda, Zhangchi Ma, Bastian Schulz, and Rune Vejlin for helpful comments and discussions. We also thank participants at seminars at Aarhus and Duke; the 1st Aarhus Workshop on Labor Markets; the Demographics, Health and Inequality (DHI) Workshop 2026; SEHO 2026; ESPE 2026; EEA–ESEM 2026; and EALE 2026. We thank Bastian Schulz for facilitating access to the Danish register data through ECONAU, Department of Economics and Business Economics, Aarhus University. Ishihata acknowledges financial support from the Tokai Foundation for Gender Studies. All errors are our own.

†Aarhus University, CESifo (almar@econ.au.dk)

‡Duke University (yusuke.ishihata@duke.edu)

1 Introduction

Fertility rates have declined substantially across high-income countries over recent decades,[1] and mitigating the decline has become a central policy objective given its consequences for the labor force and social insurance systems. A large literature accordingly studies the determinants of fertility[2] and evaluates family policies, including cash transfers, childcare provision, parental leave, and the tax treatment of couples, by analyzing already-formed couples conditional on their education (e.g., Bick, 2016; Adda et al., 2017; Jakobsen et al., forthcoming; Wang, forthcoming). This approach is well suited to the intensive margin of childbearing, but it treats partnership formation and educational attainment as fixed background characteristics rather than outcomes that may respond to policy.[3]

Two features of the data suggest that these margins matter for fertility. First, in Danish administrative data for women born in 1967–69 and 1978–80, completed fertility (the average number of children women have over their childbearing years)[4] declined across cohorts, but the decline was twice as large among all women as among women in couples because partnership rates also fell. An analysis that conditions on couples would therefore capture only about half of the decline. Second, female educational attainment rose substantially. In the later cohort, fertility was positively correlated with education, and across cohorts it declined less at higher levels of education. By shifting women toward education groups with higher fertility, educational expansion therefore offset part of the fertility decline.[5]

These descriptive patterns are consistent with several mechanisms. The positive correlation between education and fertility may arise because higher education raises earnings, makes employment and childrearing easier to combine, and changes the pool of potential partners. It may instead reflect selection, if women with different fertility preferences choose different levels of education. Similarly, the decline in partnership may reflect factors specific to partnership or a lower value of children, which reduces both fertility within couples and the gains from partnership. Changes in educational composition can also reduce partnership rates when men's and women's education distributions diverge, since individuals who prefer partners with similar education then have fewer potential matches. Marriage market conditions can in turn affect fertility by changing who forms a partnership, who partners with whom, and how bargaining power is allocated within couples. Distinguishing these channels from selection

[1] For instance, from 1980 to 2020, the total fertility rate declined from 1.84 to 1.64 in the U.S.; 1.75 to 1.33 in Japan; 1.95 to 1.82 in France; 1.68 to 1.24 in Italy; and 1.63 to 1.37 in Finland (OECD, 2026).

[2] For example, a review by Bloom et al. (2024) identifies five key determinants of fertility decline: rising income; the quantity–quality trade-off; women's opportunities, technology, and family policies; uncertainty, insurance, and intergenerational transfers; and social norms.

[3] Throughout, we use household formation, marriage, matching, relationship, and partnership to include both legal marriage and cohabitation, which are not sharply distinguished institutionally or socially in Denmark. When relevant, we distinguish whether individuals form a partnership from who partners with whom.

[4] Another widely used measure is the total fertility rate: the number of children a woman would have over her childbearing years if she experienced the age-specific fertility rates of a given year. This period measure does not describe the experience of any cohort, and it changes with the timing of births even when the number of children women eventually have does not. Completed fertility instead counts the children that women in a birth cohort actually have, regardless of when the births occur, and can be linked to each woman's education and partnership history. We focus on completed fertility because it is closer in concept to the changes in fertility behavior that we study. See, for example, Bailey (2025) for a discussion of the two measures.

[5] Doepke et al. (2023) review the evidence showing that the correlation between education and fertility was negative historically but has become increasingly positive in recent decades.

matters for policy because women induced to acquire more education or form a partnership may have different fertility preferences from those who would do so in the absence of the policy. Separating these mechanisms therefore requires modeling education, partnership formation, partner choice, bargaining, and fertility jointly and recovering the preferences behind observed choices.

This paper develops and estimates a structural equilibrium model in which individuals choose education, enter a frictionless marriage market, and then jointly choose labor supply and fertility. The marriage market clears under imperfectly transferable utility (ITU), allowing policies to affect whether and whom individuals marry and bargaining power within couples, either directly or through their effects on education. The effect of bargaining on fertility depends on how spouses value children, but couples' joint choices do not generally reveal their preferences separately. We show how variation in local marriage market conditions, together with administrative data on household choices, can separately identify women's and men's preferences for children without survey measures of desired fertility.[6] Estimated on two Danish cohorts, the model attributes the fertility decline mainly to a substantial fall in women's preferences for children, while men's preferences change little. Economic factors such as wages and the difficulty of combining work and children do not explain the decline: wage changes contribute little, and combining work and children became easier, which even offsets part of the decline. We then use the estimated model to evaluate how counterfactual policies involving childcare, joint taxation, women's wages, and education subsidies affect fertility and welfare.

Three features of the model link education and marriage market outcomes to fertility. First, the education distribution and marriage market outcomes are jointly determined in equilibrium. A policy that changes either education incentives or the gains from marriage can therefore affect all of these outcomes: the education distribution, whether and whom individuals marry, and how marital surplus is divided. Each margin can change aggregate fertility. Second, ITU makes fertility within couples responsive to bargaining power when spouses differ in their preferences for children or in the costs they bear from childbearing. Under transferable utility (TU), fertility is invariant to bargaining power, so this channel is absent.[7] Third, women differ in unobserved fertility preferences, so these preferences can shape their education choices and whether and whom they marry.

We identify women's and men's preferences for children separately using equilibrium variation in bargaining power across local marriage markets. Variation in sex ratios changes the relative scarcity of men and women and hence the Pareto weights that clear each marriage market.[8] For a given couple type, fertility choices observed in markets with different Pareto

[6] Doepke and Kindermann (2019), for example, measure disagreement using survey responses in which each spouse reports whether they want to have a child.

[7] Utility is transferable when spouses can compensate each other at a constant rate, so that what one spouse gives up equals what the other gains. A couple then chooses public goods such as children to maximize their joint surplus and uses transfers to divide it between spouses, which implies that bargaining power does not affect fertility decisions. Under imperfectly transferable utility, in contrast, the cost of a concession to one spouse is not necessarily equal to the benefit to the other, and the rate at which spouses can compensate each other varies with how resources are already divided. The couple then can no longer simply maximize joint surplus and divide it through transfers, and thus fertility decisions generally depend on the allocation of bargaining power. See Galichon et al. (2019) and Galichon and Weber (2024) for details.

[8] Pareto weights measure bargaining power within couples. We use these terms interchangeably.

weights reflect different weighted combinations of the spouses' preferences for children. Such variation allows us to distinguish women's and men's preferences. Building on Gayle and Shephard (2019), we provide a constructive identification argument for gender-specific preferences for children.

The estimated model attributes the fertility decline primarily to changes in women's preferences for children. In a structural decomposition, these changes alone would lower fertility by more than twice the model's net decline, with roughly equal contributions from lower fertility within couple types and from partnership and matching. A higher value of marriage apart from children and lower costs of combining work and children offset 149% and 31% of the net decline, respectively. Because women's preferences fell while men's did not, wives in the later cohort want fewer children than their husbands, and fertility responds more strongly to bargaining power than in the earlier cohort. The estimates provide little support for changing selection into education as an explanation for the smaller fertility decline among highly educated women.

The policy counterfactuals show that holding education and the marriage market fixed can materially understate policy effects. We implement each policy in the later-cohort economy. Our fixed-couple benchmark lets existing couples adjust labor supply and fertility while holding education, partnership, and Pareto weights at their baseline values. Free childcare raises completed fertility by 3.4% in full equilibrium, and partnership formation accounts for 41% of this gain. Low-fertility types account for a disproportionate share of the net increase in partnerships, so applying the fertility-type composition of existing partnerships would overstate the contribution of this increase to fertility. Joint taxation lowers the return to education and the value of dual-earner partnerships; both tertiary attainment and partnership formation decline. As a result, fertility falls by 2.3%, more than three times the decline in the fixed-couple benchmark.

Policies that raise women's wages can have opposite fertility effects. Removing the motherhood wage penalty raises fertility by 1.7% because it lowers the financial cost to combining work and children. Raising women's wages proportionally to close the gap between childless women and men instead lowers fertility by 2.3%, 2.8 times as much as in the fixed-couple benchmark. Higher wages improve women's opportunities outside marriage and raise their bargaining power. At the later-cohort estimates, wives prefer fewer children on average than husbands, so this bargaining response further reduces fertility.

Related Literature. This paper contributes to four strands of literature. First, we contribute to the literature on fertility decisions.[9] Within this literature, early structural models focused primarily on the trade-off between fertility and female labor supply (Caucutt et al., 2002; Francesconi, 2002; Sommer, 2016). More recent work uses richer models to evaluate policies such as cash transfers (Adda et al., 2017), childcare subsidies (Bick, 2016), parental leave (Erosa et al., 2010; Yamaguchi, 2019), tax reforms (Jakobsen et al., forthcoming), and combinations of family and labor market policies (Gayle et al., 2024; Guner et al., 2024; Kim et al., 2024; Wang,

[9] See Doepke et al. (2023), Bloom et al. (2024), Bailey (2025), and Kearney and Levine (2026) for recent comprehensive reviews of the economics literature on fertility.

forthcoming). These models generally study already-formed couples or treat partnership formation as exogenous or stochastic. Our framework evaluates family and labor market policies within an equilibrium marriage market. Policies can therefore affect aggregate fertility beyond childbearing within existing couples, through whether individuals form partnerships, who partners with whom, and how bargaining power is divided within couples. Our counterfactuals show that holding these margins fixed can substantially understate fertility responses.

Second, this paper contributes to the literature on quantitative equilibrium models of the marriage market.[10] Following Choo and Siow (2006) and the multi-market extension of Chiappori et al. (2017), this literature uses observed matching patterns to quantify the gains from marriage by gender and education. Chiappori et al. (2018) incorporate household life-cycle decisions and education choices into a TU equilibrium matching model, while Gayle and Shephard (2019) and Reynoso (2024) introduce ITU using the framework of Galichon et al. (2019). Most of these models treat fertility as exogenous or stochastic, even though children are an important part of the gains from marriage and policies studied in this literature, such as joint taxation and divorce laws, can significantly affect fertility choices. An important exception is Bisin and Tura (2026), who model endogenous fertility, cultural socialization, and separation in a TU equilibrium marriage market to study cultural integration, while taking education as given. We instead endogenize fertility within an ITU framework. Since children enter spouses' gains from marriage, fertility decisions shape the marital surplus. The marriage market, in turn, influences fertility through equilibrium Pareto weights when spouses differ in their preferences for children or the costs they bear from childbearing. Under TU, fertility decisions would not depend on spouses' relative bargaining power (Galichon et al., 2019; Galichon and Weber, 2024).

Third, our work relates to the literature on education and marriage. This literature emphasizes that returns to education arise in the marriage market as well as the labor market (Chiappori et al., 2009; Greenwood et al., 2016; Chiappori et al., 2017). Within this literature, Chiappori et al. (2018) model education choice, marriage, and labor supply but work under TU and abstract from fertility. We bring endogenous fertility into this setting. Because education is chosen before household formation and fertility, expected family outcomes enter education choices alongside labor and marriage market returns.[11] Education in turn changes the partners and household environments in which fertility choices are made. Modeling the three decisions jointly therefore captures how expected family outcomes affect education and how education changes subsequent marriage and fertility choices.

Finally, this study is among the few to model education, marriage, and fertility jointly. The closest paper is Eckstein et al. (2019), who endogenize education, marriage, employment, and fertility in a life-cycle model to account for the labor market outcomes of married women.[12] Relative to their framework, we simplify some aspects of the life cycle, such as wage dynamics, to incorporate an equilibrium marriage market with ITU. To the best of our knowledge,

[10] See Chiappori and Salanié (2023) for a recent survey of the one-to-one matching-market literature.

[11] This channel is particularly relevant given the child penalty in labor market outcomes (Kleven et al., 2019) and its variation across couple types (Almar et al., 2025).

[12] Other work models fertility and marriage together but abstracts from the equilibrium marriage market (e.g., Baudin et al., 2015; Ho and Wang, 2024).

ours is the first quantitative framework to combine an equilibrium marriage market with endogenous education and fertility choices. Three differences matter for our question: the joint determination of education and marriage, the role of bargaining, and heterogeneity in fertility preferences. In their framework, marriage offers arrive through exogenous processes, whereas our model jointly determines the education distribution and marriage market outcomes in equilibrium. Policies that change incentives for education or marriage therefore affect fertility through composition effects as well as the division of marital surplus. The models also differ in their treatment of bargaining. Their model assign spouses equal Pareto weights and identical preferences over births. We instead allow spouses to disagree over children and use an ITU framework in which fertility responds to bargaining power, the mechanism emphasized by Doepke and Kindermann (2019). Because Pareto weights are determined in the equilibrium marriage market, policies that change spouses' relative economic positions can affect fertility through bargaining. Lastly, we allow for unobserved heterogeneity in fertility preferences. This heterogeneity matters for policy because women who change their education or marriage decisions in response to a policy may have different level of fertility preferences from women who made those choices at baseline. Predicting their fertility from that of women already in the same education or partnership groups can therefore misstate the policy's effect on aggregate fertility.

The remainder of the paper is organized as follows. Section 2 describes the Danish administrative data and presents the descriptive decomposition. Section 3 presents the model of education, marriage, labor supply, and fertility. Section 4 develops the identification argument, with particular attention to spouses' preferences for children. Section 5 explains the estimation strategy, while Section 6 reports the estimates, the model's fit, and a structural decomposition of the fertility decline. Section 7 examines the policy counterfactuals. Finally, Section 8 concludes.

2 Data and Descriptives

This section describes the Danish register data and documents the cohort changes in fertility, partnership, and education that motivate the model. A descriptive decomposition shows that declines in the partnership rate and in fertility within couples make similarly sized negative contributions to the decline in cohort fertility, while the expansion of female education partly offsets them. The decomposition establishes that the partnership margin is quantitatively important; distinguishing the forces behind these changes requires the structural model.

2.1 Data

We use Danish administrative registers covering the resident population from 1980 to 2018. The registers contain information on age, birth year, legal marriage and cohabitation (BEF, Statistics Denmark, 1980–2018), education (UDDA, Statistics Denmark, 1980–2018), labor supply (RAS, Statistics Denmark, 1980–2018), and labor income (IND, Statistics Denmark, 1980–2018).

Three features of the registers are particularly useful for our analysis. First, population coverage allows us to calculate aggregate and education-specific fertility rates and the empirical distribution of educational matches without survey sampling. Second, the registers identify cohabiting as well as legally married couples. Including cohabiting couples avoids understating partnership rates and distorting measured matching patterns when unions are not legally registered. Third, each child can be linked to both parents, allowing us to construct fertility from birth records rather than self-reports.

Our primary dataset comprises all residents of Denmark observed between ages 19 and 38. We describe below the construction of the variables used in the empirical analysis and estimation.

Education. We use the Danish education register (UDDA, Statistics Denmark, 1980–2018) to measure educational attainment. We group education into four ISCED-based categories: primary education (ISCED 1–2), secondary education (ISCED 3–4), bachelor's degree (ISCED 5–6), and master's/PhD (ISCED 7–8). For each individual, educational attainment is the highest degree completed by age 38.

Marriage. The BEF register (BEF, Statistics Denmark, 1980–2018) identifies both legally married and cohabiting partners. The register classifies two individuals as cohabiting partners if they are of opposite sex, reside at the same address, differ in age by less than 15 years, and are not related. For consistency with the model, we refer to both legal marriage and cohabitation as marriage. Because the model allows each individual at most one partner, we retain the longest-lasting observed partnership, provided it lasted at least five years. We classify individuals with no qualifying partnership as single.

Labor supply. We use the RAS register (RAS, Statistics Denmark, 1980–2018) to classify each individual's labor supply. Because actual hours worked are not recorded and contracted hours are unavailable before 2008, we use the register's three categories: non-participation, part-time, and full-time.[13]

Fertility. Using parent–child links in the Danish birth register, we measure each individual's number of children. Our primary fertility measure is cohort fertility, which we define throughout the paper as the average number of children per woman accumulated between ages 19 and 38.[14] Because cohort fertility follows the same women, it can be related directly to their education and partnership histories. The total fertility rate instead combines age-specific fertility rates from different birth cohorts observed in the same calendar year and can change with

[13] These categories are imputed from discrete bins of mandatory pension contributions (ATP) paid by the employer. The threshold separating part-time from full-time employment is 30 hours per week in 1980–1992, 27 hours in 1993–2007, and 32 hours in 2008–2018 (Lund and Vejlin, 2016).

[14] Completed cohort fertility is typically measured through age 49. We use a shorter age window for two reasons. First, because the registers end in 2018, extending the age window would require drawing the later cohort from earlier birth years, thereby narrowing the gap between the two cohorts. Capping the window at age 38 allows us to compare cohorts born approximately a decade apart, so that the comparison primarily reflects across-cohort variation. Second, age-specific fertility rates are low after age 38 in cohorts observed at older ages (see Figure A.1), suggesting that this cutoff has little effect on the across-cohort comparison.

the timing of births even when completed cohort fertility is unchanged. The next subsection defines the two cohorts used in the analysis.

Local markets. We assign each individual to the administrative region in which they resided for the longest period between ages 19 and 38. In estimation, we treat each of Denmark's five administrative regions as a separate labor and marriage market: Capital Region, Region Zealand, Region Southern Denmark, Region Mid Jutland, and Region Northern Jutland. We use these regions to construct moments for local labor and marriage markets. Section 4 explains how variation across markets contributes to identification.

2.2 Patterns across cohorts

We compare women and men born in 1967–69 (the earlier cohort) with those born in 1978–80 (the later cohort).

Fertility and partnership by female education. Figure 1(a) reports cohort fertility by female education. The average number of children per woman fell by 0.043 between cohorts, with the decline concentrated among less-educated women.[15] Cohort fertility fell by 0.21 among primary-educated women and by 0.07 among secondary-educated women. It changed little among women with a bachelor's degree (-0.02) and rose slightly among women with a master's or PhD ($+0.03$).

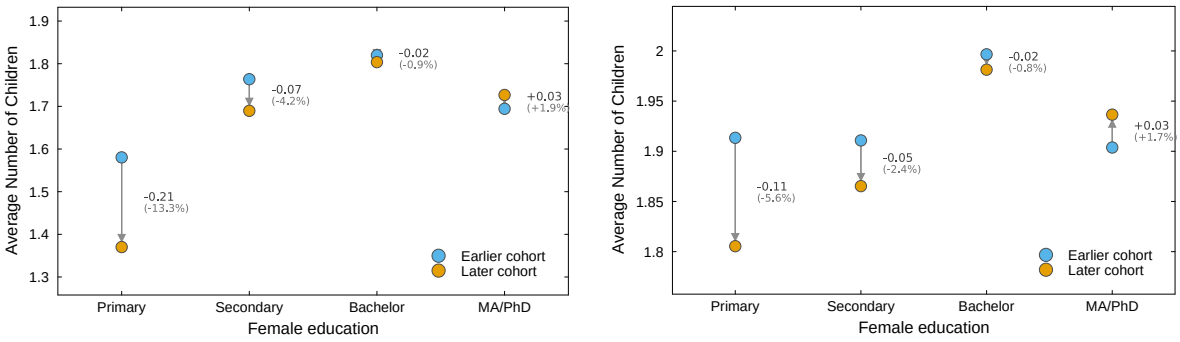
The unconditional fertility rate reflects both fertility within couples and the share of women who are in a couple. To examine fertility within couples, panel (b) of Figure 1 reports fertility conditional on being in a couple. The decline remains concentrated among less-educated women but is smaller conditional on partnership: fertility falls by 0.11 children among couples in which the woman has primary education and by 0.05 among those in which she has secondary education. It falls by 0.02 among couples in which the woman has a bachelor's degree and rises by 0.03 among those in which she has a master's or PhD. The decline among primary-educated women shrinks from -0.21 unconditionally to -0.11 conditional on partnership, suggesting that the lower partnership rate contributes meaningfully to the unconditional decline for this group.

Figure 2 reports the partnership rate by female education and cohort, using the definition in Subsection 2.1 that includes both legal marriage and cohabitation. The partnership rate fell by 6.7 percentage points (an 8.1% relative decline) among primary-educated women and by 1.7 percentage points (a 1.9% relative decline) among secondary-educated women.[16] Partnership rates were essentially unchanged among women with a bachelor's degree or a master's or

[15] The decline appears modest in part because the cohorts are only about a decade apart. A linear extrapolation over four decades yields a decline of roughly 0.17 children per woman, comparable in magnitude to declines in the total fertility rate of about 0.2 in the United States and 0.26 in Finland between 1980 and 2020 (OECD, 2026). This comparison provides a benchmark for the magnitude rather than an equivalence between the two measures, because the total fertility rate combines the behavior of multiple cohorts and can also vary with the timing of births.

[16] A partnership decline of around two percentage points appears small, but its implication for fertility is not. Couples have about 1.9 children on average. At this fertility level, a two-percentage-point decline in partnership corresponds to approximately $0.02 \times 1.9 \approx 0.04$ fewer children per woman, before accounting for changes in fertility within couples or partner education.

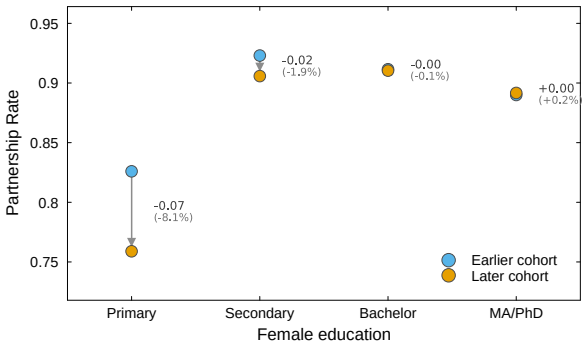
Figure 1: Average number of children per woman, by female education and cohort



(a) All women (b) Conditional on being in a couple

Notes: Data are from Danish administrative registers provided by Statistics Denmark. Women born in 1967–69 form the earlier cohort, and women born in 1978–80 form the later cohort. Cohort fertility is the average number of children per woman accumulated between ages 19 and 38. Panel (a) reports cohort fertility for all women, while panel (b) conditions on being in a couple. Each pair of points reports the earlier- and later-cohort values for an education group, with the across-cohort change reported alongside.

Figure 2: Partnership rate, by female education and cohort



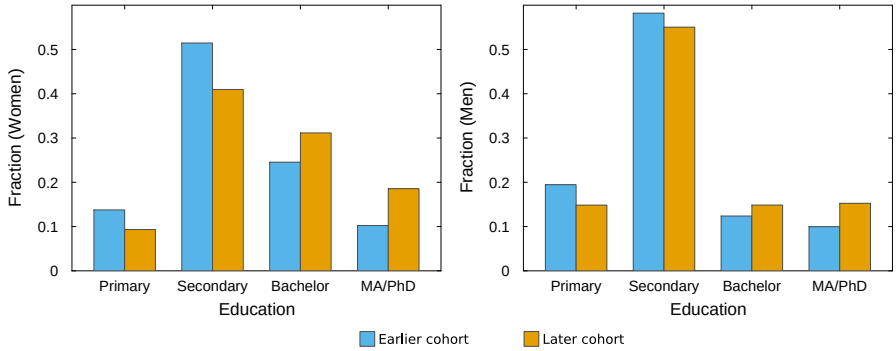
Notes: Data are from Danish administrative registers provided by Statistics Denmark. Women born in 1967–69 form the earlier cohort, and women born in 1978–80 form the later cohort. The partnership rate is the share of women classified as partnered under the definition in Subsection 2.1. The figure reports this rate by female education and cohort. Each pair of points reports the earlier- and later-cohort values for an education group, with the across-cohort change reported alongside.

PhD. The partnership-rate gradient therefore parallels the fertility gradient: both declines are concentrated among less-educated women.

Education distributions. Figure 3 reports the female and male education distributions for the two cohorts. Female educational attainment rose substantially. Among women, the share with primary education declined from 14% to 9%, and the share with secondary education declined from 51% to 41%. Over the same period, the share with a bachelor's degree increased from 25% to 31%, and the share with a master's or PhD increased from 10% to 19%. By the later cohort, 50% of women held at least a bachelor's degree, up from 35% in the earlier cohort. Men's educational attainment also rose, but by less: the share with a bachelor's degree increased from 12% to 15%, and the share with a master's or PhD increased from 10% to 15%. By the later cohort, 30% of men held at least a bachelor's degree.

This shift in the education distribution affects aggregate fertility through two channels.

Figure 3: Education distribution by gender and cohort



Notes: Data are from Danish administrative registers provided by Statistics Denmark. The earlier cohort comprises women and men born in 1967–69, and the later cohort comprises women and men born in 1978–80. Education is the highest degree completed by age 38 and is grouped into the four categories defined in Subsection 2.1. The figure reports the share of each gender-by-cohort group in each education category. Shares sum to one within gender and cohort.

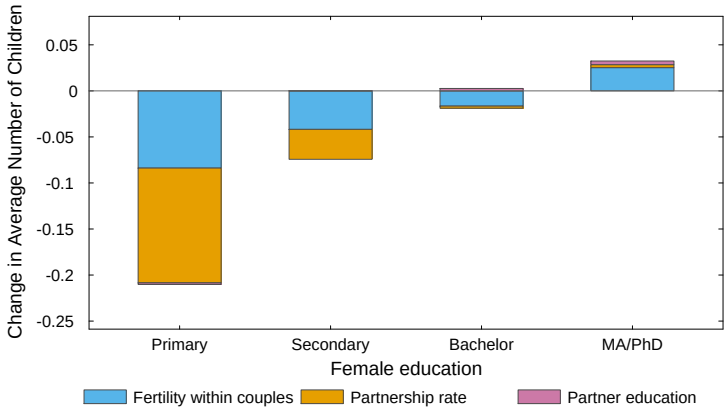
First, it mechanically reweights the population toward education groups whose fertility declined less, partially offsetting the declines within education groups described above. Second, it changes the numbers of women and men at each education level in the marriage market. Because women's education rose more than men's, the relative supplies of potential partners changed across education groups. Under assortative matching, these shifts can change both the partnership rate and the distribution of couple types, with consequences for aggregate fertility.

Descriptive decomposition. To quantify the contribution of each margin to the across-cohort change in unconditional fertility, we decompose the change into four observed components. For each female education group x, the unconditional fertility rate r_x^w can be written as the partnership rate χ_x times the matching-weighted average of fertility within couples, $\sum_y \varpi_{y|x} n_{xy}$. Here, $\varpi_{y|x}$ is the distribution of partner education conditional on a woman of education x being in a couple, and n_{xy} is the fertility rate of couple type (x, y). Aggregating across female education groups using the female education distribution W_x^{share} yields the economy-wide fertility rate R^w. We replace these components one at a time, repeat the calculation over every possible replacement order, and assign each component its average contribution across those orders. This Shapley decomposition removes dependence on any single ordering, and its contributions sum to the observed change. The partnership-rate component captures changes in whether women are partnered, while the partner-education component captures changes in the education distribution of their partners conditional on partnership. Appendix B details the construction.

Figure 4 decomposes the change within each female education group, Δr_x^w, into fertility within couples, the partnership rate, and partner education.[17] For primary-educated women, the partnership-rate contribution is -0.125, compared with -0.084 from fertility within couples. The partnership-rate and within-couple components together account for nearly all of the

[17] The female education distribution does not enter the education-specific fertility rate because conditioning on x holds female education fixed.

Figure 4: Shapley decomposition of the fertility change, by female education



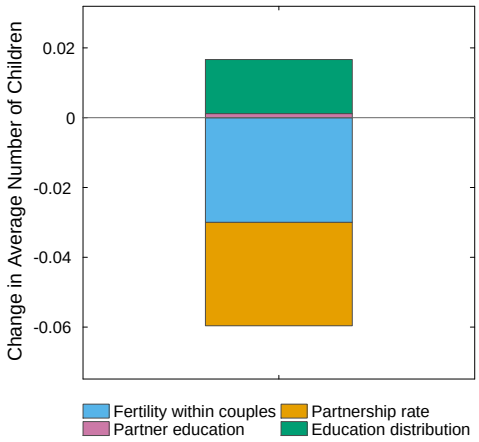
Notes: Data are from Danish administrative registers provided by Statistics Denmark. Women born in 1967–69 form the earlier cohort, and women born in 1978–80 form the later cohort. Cohort fertility is the average number of children per woman accumulated between ages 19 and 38. Within each female education group, the figure decomposes the across-cohort change into fertility within couples, the partnership rate, and partner education. Contributions are measured in children per woman. Each Shapley value averages a component's marginal contribution over all six possible replacement orders, and the contributions sum to the total change within each education group. Appendix B details the construction.

-0.210 change for this group. For secondary-educated women, fertility within couples contributes -0.042 and the partnership rate contributes -0.033. Among women with a bachelor's degree, fertility within couples contributes -0.017, while the partnership-rate and partner-education contributions are small and partly offsetting. For women with a master's or PhD, all three contributions are positive, with fertility within couples making the largest contribution, $+0.025$.

Aggregating the education-specific contributions and adding the female education distribution margin yields the economy-wide decomposition in Figure 5. Of the -0.043 aggregate decline, fertility within couples contributes -0.030, or 70% of the net decline, while the decline in the partnership rate contributes -0.030, or 69%. Changes in partner education make a small positive contribution of $+0.001$. The female education distribution contributes $+0.015$ because women shifted toward education groups with higher fertility. Educational expansion offsets 26% of the combined negative contributions of the lower partnership rate and lower fertility within couples. The aggregate decline is therefore not confined to fertility within couples: the decline in the partnership rate contributes approximately as much, while educational expansion partly offsets both.

The near-zero partner-education contribution in the female decomposition does not imply that matching patterns were unchanged. Because women's and men's education distributions changed at different rates, the distribution of men's education conditional on women's education can change little even when the distribution of women's education conditional on men's changes substantially. The corresponding male-side decompositions are in Figure A.7 and Figure A.8. The latter shows that the partner-education component offsets about 10% of the aggregate decline. This asymmetry illustrates why matching must be determined jointly from

Figure 5: Shapley decomposition of the economy-wide fertility change



Notes: Data are from Danish administrative registers provided by Statistics Denmark. Women born in 1967–69 form the earlier cohort, and women born in 1978–80 form the later cohort. Cohort fertility is the average number of children per woman accumulated between ages 19 and 38. The figure decomposes the economy-wide change into fertility within couples, the partnership rate, partner education, and the female education distribution. Contributions are measured in children per woman. Each Shapley value averages a component's marginal contribution over all 24 possible replacement orders, and the contributions sum to the economy-wide change. Appendix B details the construction.

both sides of the marriage market.

Interpretation. The decomposition is descriptive rather than causal because its four margins are jointly determined. Children contribute to the gains from partnership, so a force that lowers fertility within couples can also reduce partnership formation. Conversely, partnership formation and matching determine which couple types form and hence the specialization and bargaining conditions under which fertility is chosen. A policy can therefore move fertility within couples, partnership formation, matching, and education simultaneously.

The descriptive patterns also leave open whether women select into education based on fertility preferences. The large declines among less-educated women and roughly unchanged outcomes among highly educated women could in principle be entirely compositional. If women with stronger preferences for children increasingly select into higher education, outcomes among highly educated women may remain stable because the preference composition of this group is changing, while the lower-education groups are increasingly composed of women with weaker preferences for children. The descriptive decomposition cannot distinguish this selection-based interpretation from one in which higher education itself, through income, career flexibility, or access to potential partners, has become more conducive to childrearing across cohorts. The distinction matters for policy. If the stability of fertility among highly educated women reflects what education confers, expanding education could raise fertility. If it instead reflects selection, women induced to acquire more education need not behave like women who would have chosen that education without the policy, and expanding education need not raise fertility. To address these limitations, we develop a model that jointly determines education, marriage, and fertility, incorporates persistent heterogeneity in fertility

preferences, and allows spouses to value children differently. The model separates preference changes from selection and traces how policy affects fertility through education, the marriage market, and bargaining.

3 Model

To account for the descriptive patterns documented in Section 2.2, we develop a framework that integrates endogenous education choice, equilibrium matching in a frictionless marriage market, and subsequent intra-household decisions over labor supply and fertility.

The economy is populated by a continuum of women of mass μ_X and men of mass μ_Y. We use the labels f and m to denote women and men, respectively. Female types are indexed by x and male types by y. Within each type, individuals are indexed by i for women and j for men.

Each woman i is endowed with a fertility type $r_i \in \mathcal{R} = \{L, H\}$, representing a low or high intrinsic preference for children. Each woman knows her own fertility type, and men observe prospective partners' types in the marriage market. The econometrician does not observe fertility type. The fraction of women of fertility type r may differ across marriage markets and cohorts; we denote it by $\pi_{r,k,c}$ in market k and cohort c. When describing a single market and cohort, we suppress k and c, so the mass of type-r women is $\mu_X \pi_r$. Men are assumed to be homogeneous with respect to fertility preferences.

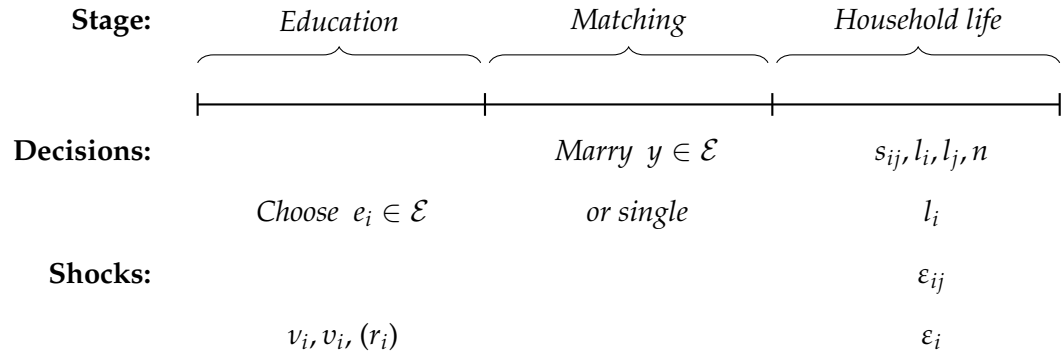
The model proceeds in three stages: education choice, matching, and household life. In the first stage, individuals choose an education level, taking into account direct education costs and idiosyncratic shocks, expected returns in the labor and marriage markets, and anticipated fertility outcomes. In the second stage, individuals participate in a one-shot, frictionless marriage market and decide whether to marry and, if so, whom to marry (Choo and Siow, 2006; Galichon et al., 2019). In the third stage, individuals enter household life as either single or married and solve a static household decision problem. Single individuals choose their labor supply, while married couples jointly choose the allocation of private consumption, each spouse's labor supply, and fertility, with potentially unequal bargaining power. Education is fixed after the first stage, and individuals do not re-enter the marriage market after the second stage. A marriage may nevertheless dissolve during household life with an exogenous probability that depends on the couple's type and the number of children. Figure 6 summarizes the timing of the model, using a woman as an illustrative example.

Agents form rational expectations about subsequent choices and equilibrium outcomes. We therefore present the model in reverse order, beginning with household life and then turning to matching and education.

3.1 Stage 3: Household Life

After choosing education and an initial marital status, individuals enter the household-life stage. We use a static discrete-choice framework in which agents choose from a finite set of alternatives, each with an additive alternative-specific shock. Because divorce can occur, the value of marriage includes the probability-weighted value of subsequent singlehood. The

Figure 6: The life cycle of a woman i



Notes: The figure summarizes the timing of decisions and shocks for a woman i. In the first stage, she chooses education e_i. In the second stage, she chooses a partner's education y or remains single. In the third stage, a single woman chooses labor supply l_i, while a married couple chooses the woman's share of private consumption s_{ij}, the spouses' labor supplies (l_i, l_j), and fertility n. The variables v_i and v_i denote the matching and education-cost shocks, respectively; r_i denotes the woman's fertility type and is realized before any decisions are made; and ε_i and ε_{ij} denote the alternative-specific shocks for singles and couples. Men face the analogous timing.

value of singlehood does not include future marriage because remarriage is ruled out. We first characterize singlehood and then use its value to construct the value of marriage.

3.1.1 Single's problem

A single woman i is characterized by her type $x = (e, r)$, where e denotes her education and r indicates her fertility type. She chooses her labor supply $l_i \in \{0, 0.75, 1\}$, corresponding to not working, part-time, and full-time employment, respectively.[18]

New births occur only within marriage in the model. A single woman may nevertheless have children from a previous marriage if she is divorced. We assume that children reside with their mother following divorce, so a divorced mother's state includes the number of children $n \in \{0, 1, 2, 3\}$.[19] We do not otherwise distinguish between never-married and divorced women: a divorced woman with no children has the same decision problem as a never-married woman.

Flow utility. Let $u_f^{x0}(l_i; n)$ denote the utility of a single woman i of type x, given her labor supply choice l_i and her number of children n.[20] Her flow utility is specified as follows:

$$\begin{aligned} u_f^{x0}(l_i; n) = \ln q_i &- \beta_f \left[\mathbb{I}(l_i = 1) + \beta_f^{PT} \mathbb{I}(l_i = 0.75) \right] \\ &+ \Gamma_{f,r}^{s}(n) \\ &- \alpha_f^{s}(n; e) \left[\mathbb{I}(l_i = 1) + \alpha_f^{PT} \mathbb{I}(l_i = 0.75) \right] \end{aligned} \tag{1}$$

[18] Following Jakobsen et al. (forthcoming), we assume that part-time employment corresponds to 75% of full-time hours. This assumption reflects institutional features of the Danish labor market, where standard full-time work is 37 hours per week and part-time schedules typically range between 28 and 32 hours per week.

[19] We top-code the number of children at three, so $n = 3$ represents three or more children.

[20] The notation u_f^{x0} represents the utility of a woman of type x when she is single, that is, when her partner's type is null, denoted by 0. A never-married woman has $n = 0$.

where her consumption q_i is determined by the following budget constraint:

$$q_i = (1 - \kappa(n))(w_i l_i + \mathbb{I}(l_i = 0)\mathcal{B} - \mathcal{T}(w_i, l_i, n) - C^S(w_i l_i, n) + \mathcal{D}(n)). \tag{2}$$

The first term in Equation (1) represents utility from private consumption. The second term captures the disutility of labor: β_f is the disutility of full-time work, part-time work ($l_i = 0.75$) incurs the fraction $\beta_f^{PT} \in (0,1)$ of this disutility, and not working ($l_i = 0$) carries none. We use the same specification for men, replacing β_f with β_m. The third term, $\Gamma^s_{f,r}(n)$, is the utility a divorced woman derives from the children of her previous marriage and is allowed to vary with her fertility type r. The final term captures the cost of balancing work and family responsibilities and is governed by $\alpha^s_f(n;e)$. This function is the single-mother counterpart of the partnered cost $\alpha_f(n;e)$, specified in the couple's problem below. We allow the cost of combining work and childrearing to differ between single and married mothers. For single mothers, the estimated specification has no fixed component: $\alpha^s_f(n;e) = \alpha^s_{1,f,e} n + \alpha^s_{2,f} n^2$. As with the disutility of labor, part-time work incurs a fraction $\alpha_f^{PT} \in (0,1)$ of the work–family cost incurred under full-time work.

In the budget constraint (Equation (2)), a working woman earns $w_i l_i$, while a non-employed woman receives benefit $\mathcal{B}$; both are net of taxes $\mathcal{T}(w_i, l_i, n)$. A divorced mother also pays net childcare costs $C^S(w_i l_i, n)$ and receives child support $\mathcal{D}(n)$. The tax, childcare, benefit, and child-support schedules are calibrated to the Danish policy environment, as described in Appendix C. Finally, the factor $1 - \kappa(n)$ converts household resources into private consumption after accounting for expenditures on children.[21]

Value functions. Let $\varepsilon_i(l_i)$ be the alternative-specific shock term. The solution to the single woman's problem is

$$l_i^*(n; \varepsilon_i(l_i)) = \arg\max_{l_i} \left\{ u_f^{x0}(l_i; n) + \varepsilon_i(l_i) \right\}. \tag{3}$$

Assuming the alternative-specific shocks ε_i follow a Type-I extreme value distribution with scale parameter σ_ε, we obtain a closed-form expression for the expected value of being single, EV_f^{x0}:

$$\begin{aligned} EV_f^{x0}(n) &= E \max_l \left[u_f^{x0}(l; n) + \varepsilon_i(l) \right] \\ &= \sigma_\varepsilon c + \sigma_\varepsilon \ln \left[\sum_l \exp\left(u_f^{x0}(l; n)/\sigma_\varepsilon \right) \right] \end{aligned} \tag{4}$$

where c is the Euler–Mascheroni constant.

The conditional choice probability of selecting labor supply l given children n is then

$$\mathbb{P}_f^{x0}(l; n) = \frac{\exp\left(u_f^{x0}(l; n)/\sigma_\varepsilon \right)}{\sum_{l'} \exp\left(u_f^{x0}(l'; n)/\sigma_\varepsilon \right)}. \tag{5}$$

[21] We use the square root equivalence scale: $\kappa(n) = 1 - \sqrt{2/(2+n)}$ (Eckstein et al., 2019).

Single man's problem. A single man j of type y chooses his labor supply $l_j \in \{0, 0.75, 1\}$. As with single women, his number of children $n \in \{0, 1, 2, 3\}$ records children from a previous marriage, and a never-married man has $n = 0$. Children do not enter the single man's flow utility directly; instead, n affects his decision problem through the budget constraint, via child support payments $\mathcal{D}(n)$. His flow utility is therefore

$$u_m^{0y}(l_j; n) = \ln q_j - \beta_m \Big[\mathbb{I}(l_j = 1) + \beta_m^{PT}\, \mathbb{I}(l_j = 0.75)\Big] \tag{6}$$

where his consumption q_j is determined by the budget constraint

$$q_j = w_j l_j + \mathbb{I}(l_j = 0)\mathcal{B} - \mathcal{T}(w_j, l_j, 0) - \mathcal{D}(n). \tag{7}$$

Taxes are evaluated at zero children because child-related transfers accrue to the custodial parent, whom we take to be the mother. The expected value of being single $EV_m^{0y}(n)$ and the conditional choice probabilities $\mathbb{P}_m^{0y}(l; n)$ are defined analogously to the single woman's case, with shocks $\varepsilon_j(l_j)$ drawn from a Type-I extreme value distribution with the same scale parameter σ_ε.

3.1.2 Couple's problem

Consider a couple consisting of woman i of type x and man j of type y. The couple jointly chooses (i) the woman's share of private consumption, $s_{ij} \in [0, 1]$ (with the man receiving $1 - s_{ij}$), (ii) labor supply for each spouse, $l_i, l_j \in \{0, 0.75, 1\}$, and (iii) the number of children, $n \in \{0, 1, 2, 3\}$. For notational convenience, let $a = (s_{ij}, l_i, l_j, n) \in \mathcal{A}^M$ denote the vector of decision variables, where $\mathcal{A}^M$ represents the set of possible choice alternatives.

Flow utility. Let $u_f^{xy}(a)$ and $u_m^{xy}(a)$ denote the period utility of a married woman i and a married man j, respectively, in a type-xy couple. These are specified as follows:

$$\begin{aligned} u_f^{xy}(a) = \ln(\psi s_{ij} q) &- \beta_f \Big[\mathbb{I}(l_i = 1) + \beta_f^{PT}\, \mathbb{I}(l_i = 0.75)\Big] \\ &+ \Gamma_{f,r}(n) \\ &- \alpha_f(n; e) \Big[\mathbb{I}(l_i = 1) + \alpha_f^{PT}\, \mathbb{I}(l_i = 0.75)\Big] \end{aligned} \tag{8}$$

$$\begin{aligned} u_m^{xy}(a) = \ln(\psi(1 - s_{ij}) q) &- \beta_m \Big[\mathbb{I}(l_j = 1) + \beta_m^{PT}\, \mathbb{I}(l_j = 0.75)\Big] \\ &+ \Gamma_m(n) \\ &- \alpha_m(n; y) \Big[\mathbb{I}(l_j = 1) + \alpha_m^{PT}\, \mathbb{I}(l_j = 0.75)\Big] \end{aligned} \tag{9}$$

where household consumption q is given by the budget constraint

$$q = (1 - \kappa(n))(w_i l_i + w_j l_j + \mathbb{I}(l_i = 0)\mathcal{B} + \mathbb{I}(l_j = 0)\mathcal{B} - \mathcal{T}(w_i, w_j, l_i, l_j, n) - C(w_i, w_j, l_i, l_j, n)). \tag{10}$$

The parameter ψ captures economies of scale from sharing household consumption between two adults. We set $\psi = \sqrt{2}$, consistent with the square-root equivalence scale (Eckstein et al., 2019), so a childless couple requires $\sqrt{2}$ times the resources of one single adult, rather than twice those resources, to give both spouses the same consumption utility under equal sharing.[22]

The functions $\Gamma_{g,r}$ and α_g capture, respectively, the utility from children and the cost of combining work with children. We use parsimonious quadratic specifications that allow the marginal value and work–family cost of an additional child to vary with the number of children:

$$\Gamma_{g,r}(n) = \gamma_0 \mathbf{1}\{n \geq 1\} + \gamma_{1,g,r}\, n + \gamma_2\, n^2, \tag{11}$$

$$\alpha_g(n; e) = \alpha_{0,g} \mathbf{1}\{n \geq 1\} + \alpha_{1,g,e}\, n + \alpha_{2,g}\, n^2, \tag{12}$$

both zero at $n = 0$ by construction. Conditional on gender and, for women, fertility type, intrinsic utility from children does not depend directly on either spouse's education. Partner type instead affects fertility through household income, work–family costs, dissolution risk, and the equilibrium Pareto weight.

In (11), $\gamma_{1,g,r}$ is gender-specific and, for women, fertility-type-specific. The restriction $\gamma_{1,f,H} > \gamma_{1,f,L}$ ensures that high-type women derive greater utility from children, while men share the common coefficient $\gamma_{1,m}$. The common intercept γ_0 separates the value of becoming a parent from the value of additional children, while γ_2 governs how marginal utility changes with the number of children. A negative γ_2, for instance, makes the marginal utility of another child decline with family size.

In (12), $\alpha_{0,g}$ is a fixed work–family cost incurred when a working parent has any children. The linear coefficient $\alpha_{1,g,e}$ allows the additional cost per child to vary by education, while $\alpha_{2,g}$ allows that cost to change with the number of children and is common across education within gender. Labor supply enters through the intensity indicator: a parent working part-time bears a fraction $\alpha_g^{PT} \in (0, 1)$ of the cost borne at full-time work. The work–family term for an employed parent is therefore $\alpha_g(n; e)\left[\mathbb{I}(l = 1) + \alpha_g^{PT}\, \mathbb{I}(l = 0.75)\right]$, and it enters utility with a negative sign. We do not restrict the sign of α_g: a negative value represents a pull toward employment when children are present rather than a cost of it.

The household's period utility is a weighted average of the spouses' utilities:

$$u^{xy}(a; \lambda^{xy}) = \lambda^{xy} u_f^{xy}(a) + (1 - \lambda^{xy}) u_m^{xy}(a) \tag{13}$$

where $\lambda^{xy} \in (0, 1)$ is the Pareto weight assigned to the female spouse.

This specification allows for a closed-form solution of the consumption share s_{ij}, which depends solely on the Pareto weight λ^{xy}.[23] The first-order condition for the consumption share

[22] Under log utility ψ enters each spouse's utility additively as $\ln \psi$, so it does not affect the couple's labor-supply and fertility choices, which depend only on utility differences across alternatives. It does shift the level of the value of marriage relative to singlehood, and therefore continues to affect the matching and education decisions and the equilibrium Pareto weights.

[23] The marriage market equilibrium determines the Pareto weight for each couple type and market. Because household life is static, the weight and the resulting consumption share do not vary within that stage.

yields

$$s_{ij}(\lambda^{xy}) = \lambda^{xy}, \tag{14}$$

that is, under log utility the female's share of private consumption equals her Pareto weight, and is therefore strictly increasing in λ^{xy}. Because the consumption share s_{ij} is uniquely determined by the equilibrium Pareto weight, we exclude it from the set of choice variables in the analysis that follows. Hence, the vector of decision variables for a married couple is given by $a = (l_i, l_j, n)$, and we henceforth let $\mathcal{A}^M$ denote this reduced set of alternatives. We correspondingly write the spouses' flow utilities as $u_f^{xy}(a; \lambda^{xy})$ and $u_m^{xy}(a; \lambda^{xy})$, making explicit that they depend on the Pareto weight through the consumption share.

Value functions. Let $V_f^{xy}(a; \lambda^{xy})$ denote the choice-specific value function for woman i in a type-xy couple. The term $\varepsilon_{ij}(a)$ is a household-level shock to the joint allocation a and therefore enters both spouses' evaluations of that allocation. The total value associated with action $a = (l_i, l_j, n)$ is given by:

$$\begin{aligned} &V_f^{xy}(a; \lambda^{xy}) + \varepsilon_{ij}(a) \\ &= (1 - d^{xy}(n)) u_f^{xy}(a; \lambda^{xy}) + \varepsilon_{ij}(a) + d^{xy}(n) E V_f^{x0}(n) \end{aligned} \tag{15}$$

where $d^{xy}(n)$ is the exogenous divorce probability, which depends on the couple type and the number of children. This formulation of the value functions embodies two assumptions. First, although household decisions are static, we allow for the possibility of divorce. Equivalently, $d^{xy}(n)$ can be read as the average fraction of the (post-childbearing) life cycle a type-xy marriage spends dissolved, so that $1 - d^{xy}(n)$ weights the married payoff and $d^{xy}(n)$ the divorced payoff in the value function. Under this reading, fertility outcomes are evaluated over the full life cycle, while labor supply can be adjusted following divorce, implying an implicit within-period timing in which fertility decisions precede potential marital dissolution. This structure allows us to examine the role of changes in divorce risk in shaping fertility outcomes, even within a static decision framework. Second, divorce probabilities are allowed to vary with partner type and the number of children. Because the dissolution probability varies with the number of children, observed differences in dissolution risk enter the couple's fertility incentives. In particular, a lower value of $d^{xy}(n)$ raises the expected weight placed on the married payoff.

The household's choice-specific value function (plus shock) for a type-xy couple is given by

$$V^{xy}(a; \lambda^{xy}) + \varepsilon_{ij}(a) = \lambda^{xy} V_f^{xy}(a; \lambda^{xy}) + (1 - \lambda^{xy}) V_m^{xy}(a; \lambda^{xy}) + \varepsilon_{ij}(a) \tag{16}$$

where $V_m^{xy}(a; \lambda^{xy})$ is the male counterpart to the female value function. The household's optimal choice is then determined by

$$a_{ij}^*(\lambda^{xy}, \varepsilon_{ij}(a)) = \arg\max_a \left\{ V^{xy}(a; \lambda^{xy}) + \varepsilon_{ij}(a) \right\} \tag{17}$$

We assume that the household's alternative-specific utility shocks, $\varepsilon_{ij}(a)$, follow a nested logit structure with scale parameter σ_ε (shared with the single's problem) and nesting parameter ρ_ε. Alternatives are grouped by fertility, so labor-supply combinations associated with the same number of children may have correlated unobserved values and be closer substitutes than alternatives with different fertility levels. Under this assumption, the expected value function $EV_f^{xy}(\lambda^{xy})$ admits a closed-form expression:

$$\begin{aligned} EV_f^{xy}(\lambda^{xy}) &= E\left[V_f^{xy}(a^*(\lambda^{xy});\lambda^{xy}) + \varepsilon_{ij}(a^*)|\lambda^{xy}\right] \\ &= \sigma_\varepsilon c + \sigma_\varepsilon \ln\left(\sum_{n'=0}^{3} S_{n'}^{\rho_\varepsilon}(\lambda^{xy})\right) + \sum_a \mathbb{P}^{xy}(a;\lambda^{xy})\left[V_f^{xy}(a;\lambda^{xy}) - V^{xy}(a;\lambda^{xy})\right] \end{aligned} \tag{18}$$

where c is the Euler–Mascheroni constant,

$$S_n(\lambda^{xy}) = \sum_{l_i,l_j} \exp\left(V^{xy}((l_i,l_j,n);\lambda^{xy})/\sigma_\varepsilon\rho_\varepsilon\right) \tag{19}$$

is the inclusive value of the nest in which the couple has n children, and

$$\mathbb{P}^{xy}(a=(l_i,l_j,n);\lambda^{xy}) = \frac{\exp\left(V^{xy}(a;\lambda^{xy})/\sigma_\varepsilon\rho_\varepsilon\right)}{S_n(\lambda^{xy})} \frac{S_n^{\rho_\varepsilon}(\lambda^{xy})}{\sum_{n'=0}^{3} S_{n'}^{\rho_\varepsilon}(\lambda^{xy})} \tag{20}$$

is the conditional choice probability of the alternative a for a type-xy couple with the Pareto weight λ^{xy}.

3.1.3 Wages

We assume that the wage depends on gender g, education e, and the number of children n. We write this wage as $w_{g,e,n}$, so that a woman of education e with n children earns $w_i = w_{f,e,n}$ and a man of education y earns $w_j = w_{m,y,n}$. Letting the wage vary with gender and education captures the market return to schooling and the gender wage gap, both of which shape the value of each education level and of each marital match. Female wages vary with the number of children to capture the observed motherhood wage penalty in reduced form (Kleven et al., 2019).[24]

3.2 Stage 2: Matching

In the *Matching* stage, agents participate in a one-time, frictionless marriage market and decide whether, and whom, to marry, as in the separable matching framework of Choo and Siow (2006) and its imperfectly transferable extension in Galichon et al. (2019). This process results in $|\mathcal{R}||\mathcal{E}|^2 + |\mathcal{R}||\mathcal{E}| + |\mathcal{E}|$ possible household types: $|\mathcal{R}||\mathcal{E}|^2$ types of married couples, $|\mathcal{R}||\mathcal{E}|$ types of single women, and $|\mathcal{E}|$ types of single men, where $|\mathcal{R}|$ is the number of fertility types and $|\mathcal{E}|$ is the number of education categories.

For a woman i of type x, the total value of matching with a man j of type y is denoted by

[24] Male wages are assumed to be independent of the number of children, that is, $w_{m,e,n} = w_{m,e}$. This assumption is also consistent with the findings of Kleven et al. (2019), who show that men do not experience child penalties.

$U^{xy}_{f,ij}(\lambda^{xy})$, where λ^{xy} is the Pareto weight assigned to the woman in a type-xy couple. This total value comprises three components:

$$U^{xy}_{f,ij}(\lambda^{xy}) = EV^{xy}_{f}(\lambda^{xy}) + \xi^{xy} + \nu^{xy}_{ij} \tag{21}$$

where ν^{xy}_{ij} is an idiosyncratic taste shock specific to a man j of type y (Gutierrez, 2020). The total matching value for a man j of type y with a woman i of type x, $U^{xy}_{m,ji}(\lambda^{xy})$, is defined analogously.

The first term, $EV^{xy}_{f}(\lambda^{xy})$, represents the economic value common to all women of type x who match with men of type y. This economic component reflects the expected value of entering a type-xy marriage and is endogenously determined by intra-household allocation decisions during the household-life stage. This term also depends on equilibrium conditions in the marriage market, specifically the demand and supply for each type, since the Pareto weight λ^{xy} is determined as part of the equilibrium and serves as a "price" that clears the market. For instance, if there is excess demand for women of x from men of type y, men of type y must transfer part of their utility by increasing λ^{xy}, thereby increasing $EV^{xy}_{f}(\lambda^{xy})$; see the solution algorithm described in Appendix D for details.

The second term, ξ^{xy}, captures match quality for each couple type, interpreted as non-economic motives for forming a particular partnership.[25] Match quality provides an additional force toward assortative matching and has been shown to matter for fitting observed matching patterns in the data (Chiappori et al., 2018; Eckstein et al., 2019). We parameterize match quality using the direction and size of the spouses' education gap:

$$\xi^{xy} = \xi_0 + \xi_1 \max(y - e, 0) + \xi_2 \max(e - y, 0). \tag{22}$$

The intercept ξ_0 is the non-economic value of marriage relative to singlehood. The coefficient ξ_1 applies when the husband is more educated and ξ_2 when the wife is more educated, with each term increasing linearly in the size of the gap. This parsimonious specification allows non-economic match quality to vary with both the size and direction of the spouses' education gap. Conditional on the size and direction of this gap, the remaining systematic variation across couple types comes from the model's economic value of marriage.[26]

After drawing the vector of idiosyncratic tastes $\nu^{x}_{i} = \{\nu^{xy}_{ij}\}_{y\in\{0\}\cup\mathcal{E},j}$, where $y = 0$ denotes the option of remaining single, a woman of type x chooses her preferred partner type (or singlehood) to maximize her total marital value:

$$y^{*} = \underset{y\in\{0\}\cup\mathcal{E}}{\arg\max} \left\{ EV^{xy}_{f}(\lambda^{xy}) + \xi^{xy} + \max_{j} \nu^{xy}_{ij} \right\}. \tag{23}$$

[25] This term also absorbs any mechanical, alternative-count-driven utility difference between married and single, which would otherwise be confounded with non-economic match value. Because married couples choose over a larger alternative set than singles (the fertility dimension is added), the expected maximum of their utility shock draws is mechanically larger; we let ξ^{xy} subsume this so that it does not load onto the economic component EV^{xy}_{f}.

[26] We have $\xi^{xy} = 0$ for singles.

The utility of a woman i of type x with marital preferences ν_i^x is then defined as

$$A_F(x) = \max_{y \in \{0\} \cup \mathcal{E}} \left\{ EV_f^{xy}(\lambda^{xy}) + \xi^{xy} + \max_j \nu_{ij}^{xy} \right\} \tag{24}$$

where the dependence on her own draws ν_i^x is suppressed in the notation. The counterpart for a man of type y, $A_M(y)$, is defined analogously.

3.3 Stage 1: Education Choice

In the first stage, each individual chooses an education level e from the set $\mathcal{E}$. For each gender g, completing education level e entails a cost $c_{g,e}(X_i, v_{ei}) = c_{g,e}(X_i) + v_{ei}$, where X_i represents observed characteristics (e.g., parental income), and v_{ei} captures unobserved costs, such as psychological or cognitive burdens. Following Chiappori et al. (2018), individuals observe their idiosyncratic preferences over potential partners before choosing education. This implies that they know the utility of marriage with each partner type, as well as the utility of singlehood, as given in Equation (24). A woman i of fertility type r selects the education level e^* that maximizes her expected utility in subsequent periods, that is,

$$e^* = \arg\max_{e \in \mathcal{E}} \left\{ A_F(e, r) - c_{f,e}(X_i, v_{ei}) \right\} \tag{25}$$

where $A_F(x)$ is defined in Equation (24) and $x = (e, r)$ denotes the woman's type. Similarly, a man j chooses his education level y^* to maximize his expected utility in subsequent periods:

$$y^* = \arg\max_{y \in \mathcal{E}} \left\{ A_M(y) - c_{m,y}(X_j, v_{yj}) \right\}. \tag{26}$$

Individuals, therefore, consider labor market returns, marriage market prospects, and fertility-related trade-offs when making their education decisions.

Because the gains from education depend on the education choices made on both sides of the marriage market, the original sequential education game may have multiple equilibria. Following Nöldeke and Samuelson (2015) and Chiappori et al. (2018), we select the equilibrium induced by an auxiliary game in which education and matching are contracted jointly before education is sunk. Agents observe their education-cost and marital-preference shocks when participating in this market, and a prospective pair can condition its match on both partners' education choices. In the competitive representation of this auxiliary game, agents take the complete schedules of Pareto weights and type masses as given and choose their own education and partner type. Nöldeke and Samuelson (2015) characterize an ex-ante equilibrium of this kind as pairwise efficient and show that it is also an ex-post equilibrium of the original sequential game; existence of the auxiliary equilibrium in our environment is established in Proposition 1. Accordingly, a woman i of fertility type r jointly chooses her education level e and a partner type y (including the option of remaining single, $y = 0$), conditional on her

idiosyncratic preferences ν_{ij}^{xy} and education cost $c_{f,e}(X_i, v_{ei}) = c_{f,e}(X_i) + v_{ei}$:

$$(e^*, y^*) = \underset{(e,y)\in\mathcal{E}\times(\{0\}\cup\mathcal{E})}{\arg\max} \left\{ EV_f^{xy}(\lambda^{xy}) + \xi^{xy} + \max_j \nu_{ij}^{xy} - (c_{f,e}(X_i) + v_{ei}) \right\} \tag{27}$$

where $x = (e, r)$ denotes the woman's type.

We assume that the shock to the matching decision ν_{ij}^{xy} follows a Type-I extreme value distribution with scale parameter σ_ν, while the shock to education costs v_{ei} follows a normal distribution with mean zero and standard deviation σ_{ec}. We further assume that the deterministic component of the education cost is homogeneous conditional on the education level, $c_{f,e}(X_i) = c_{f,e}$. Conditional on the realization of v_e, we then obtain a closed-form expression for the probability that a woman of fertility type r chooses education e^*, so that her type is $x^* = (e^*, r)$, and matches with a man of education y^* (Gutierrez, 2020):

$$\mu_f^{x^*y^*}(\lambda) = \frac{m_{y^*} \exp\left(\left[EV_f^{x^*y^*}(\lambda^{x^*y^*}) + \xi^{x^*y^*} - c_{f,e^*} - v_{e^*} \right] / \sigma_\nu \right)}{\sum_{e\in\mathcal{E}} \sum_{y\in\{0\}\cup\mathcal{E}} m_y \exp\left(\left[EV_f^{xy}(\lambda^{xy}) + \xi^{xy} - c_{f,e} - v_e \right] / \sigma_\nu \right)} \tag{28}$$

where the woman's fertility type r is held fixed in the denominator, with $x = (e, r)$, since it is an endowment rather than a choice. Because a woman's type is $x = (e, r)$ whereas a man's type is his education y, men choose among prospective partners who differ in both education and fertility type, while women choose among men who differ only in education. The analogous probability that a man chooses education level y^* and matches with a woman of type x^* is

$$\mu_m^{x^*y^*}(\lambda) = \frac{f_{x^*} \exp\left(\left[EV_m^{x^*y^*}(\lambda^{x^*y^*}) + \xi^{x^*y^*} - c_{m,y^*} - v_{y^*} \right] / \sigma_\nu \right)}{\sum_{y\in\mathcal{E}} \sum_{x\in\{0\}\cup(\mathcal{E}\times\mathcal{R})} f_x \exp\left(\left[EV_m^{xy}(\lambda^{xy}) + \xi^{xy} - c_{m,y} - v_y \right] / \sigma_\nu \right)} \tag{29}$$

where f_x and m_y denote measures of women and men who choose type x and y, respectively, with the convention $m_0 = f_0 = 1$ for the singlehood option. The mass multiplier arises because the relevant taste draw is the maximum of ν over the mass of potential partners of each type.

Because the shocks are match-specific, the probability of choosing a partner type depends explicitly on the mass of potential partners of that type (Gutierrez, 2020). Consequently, splitting a partner type into otherwise identical subtypes simply divides its mass between them and leaves their combined matching probability unchanged. A standard type-level multinomial logit instead treats the new labels as additional alternatives, so such a split mechanically changes matching probabilities and the Pareto weights required to clear the market. This invariance is particularly valuable in our setting because female education categories are further divided by latent fertility type.

We integrate these probabilities over the education-cost shocks. This allows substitution across education choices to depart from multinomial logit, while partner-type choices conditional on education retain the logit form (Chiappori et al., 2018).

3.4 Equilibrium

We solve for a competitive auxiliary equilibrium in which households choose optimal allocations and agents choose education and partner types optimally in the auxiliary market. Single and married individuals first solve their household-life problems; the resulting values enter the joint education-and-partner choices. We formally define the equilibrium in Definition 1.

Definition 1 (Auxiliary equilibrium). *A competitive auxiliary equilibrium consists of (i) labor supply choices for single women $l_i^*(n;\varepsilon_i)$ and single men $l_j^*(n;\varepsilon_j)$, and allocation choices for married couples $a_{ij}^*(\lambda,\varepsilon_{ij})$, (ii) a matrix of the Pareto weights $\{\lambda^{xy,*}\}_{x\in\mathcal{E}\times\mathcal{R},y\in\mathcal{E}}$, (iii) choice probabilities of education and partner type for women $\mu_f^{xy}(\lambda)$ and men $\mu_m^{xy}(\lambda)$, integrated over the education-cost shocks, and (iv) type masses $\{f_x, m_y\}$ such that*

1. *Household allocation decisions solve Equations (3) and (17).*
2. *Education and partner-type choices solve the joint choice problems that yield Equations (28) and (29).*
3. *For all $x=(e,r)\in\mathcal{E}\times\mathcal{R}, y\in\mathcal{E}$, the measure of women who choose type x as their own type and type y as their partner type equals the measure of men who choose type y as their own type and type x as their partner type:*

$$\mu_{\mathcal{X}}\pi_r\mu_f^{xy}(\lambda^*)=\mu_{\mathcal{Y}}\mu_m^{xy}(\lambda^*) \tag{30}$$

   where the choice probabilities are given by Equations (28) and (29).
4. *The type masses are consistent with the choices that generate them: $f_x=\mu_{\mathcal{X}}\pi_r\sum_{y\in\{0\}\cup\mathcal{E}}\mu_f^{xy}(\lambda^*)$ for all $x=(e,r)$ and $m_y=\mu_{\mathcal{Y}}\sum_{x\in\{0\}\cup(\mathcal{E}\times\mathcal{R})}\mu_m^{xy}(\lambda^*)$ for all y.*

The following proposition establishes that an equilibrium exists.

Proposition 1. *Suppose the matching taste shocks ν follow the Type-I extreme-value distribution and are independent of the almost-surely finite education-cost shocks v. If the optimized spouse values satisfy the sufficient conditions stated in Appendix E, then a competitive auxiliary equilibrium exists, and its Pareto weights satisfy $\lambda^{xy}\in(0,1)$ for every couple type (x,y).*

Proof of Proposition 1. See Appendix E. □

The proof extends the excess-demand construction of Gayle and Shephard (2019) to a single fixed point over the Pareto weights and endogenous education-type masses. Following Chiappori et al. (2018), we use the auxiliary-game allocation as our equilibrium selection; they argue that this allocation is generically unique, although uniqueness is not part of Proposition 1. We explain the numerical solution algorithm in Appendix D.

3.5 Key Mechanisms

Five mechanisms link education and marriage to fertility: the value and cost of children, bargaining within couples, partnership formation, partner composition, and selection by fertility type.

First, fertility within a couple depends on the value each spouse places on children and the costs of raising them. The functions $\Gamma_{f,r}$ and Γ_m allow the marginal value of children to vary with family size and, for women, fertility type. Children reduce private consumption through child-related expenditure and childcare costs; for working parents, they also interact with work–family costs, while the motherhood wage penalty lowers women's earnings when children are present (Kleven et al., 2019). Family-size choices reflect both valuations and costs, while labor-supply choices provide additional information about the costs of combining employment and children.

Second, bargaining affects fertility when spouses differ in their utility from children, the costs they bear from childbearing, or both. Doepke and Kindermann (2019) document survey evidence that spouses disagree about having another child. Under ITU, the Pareto weight determines how spouses' net valuations enter their fertility choice; under TU, changing the division of marital surplus would not change that choice. The equilibrium Pareto weight responds to the relative supply of partner types and to spouses' outside options. Because fertility is a joint choice, observed family size alone does not reveal women's and men's preferences for children separately. Section 4 shows how variation in equilibrium Pareto weights across markets, together with household choices, identifies these preferences.

Third, aggregate fertility depends on whether couples form. The economic value of marriage and the match value ξ^{xy} determine its value relative to singlehood. Since new births occur only within partnerships in the model, changes in the gains from marriage can affect fertility even if existing couples do not change their choices. Children themselves enter those gains, so a change in their value can affect both partnership formation and fertility within couples.

Fourth, who partners with whom also matters. Couple types can differ in household income, work–family costs, dissolution risk, and equilibrium Pareto weights. Changing the distribution of matches therefore changes aggregate fertility by reweighting couple types, even if the partnership rate and fertility within each type remain fixed. Changes in the relative supply of partner types can also alter Pareto weights and, through bargaining, fertility within couples.

Finally, women's persistent fertility types enter their education and partner choices. The correlation between education and fertility can reflect both the economic consequences of education and selection by preferences for children. Women induced by a policy to acquire education or enter a partnership need not have the same fertility preferences as women who made those choices without the policy. Extrapolating from women who would make those choices anyway can therefore misstate a policy's effect on aggregate fertility.

4 Identification

Three objects central to this paper are not directly observed: how preferences for children vary across women, how those preferences differ between spouses, and the Pareto weights governing couples' choices. A couple's completed family size is a joint outcome and does not reveal these objects separately. This section establishes conditions under which auxiliary measurements, marriage market outcomes, and household choices identify them.

The identification argument proceeds in three steps. Subsection 4.1 recovers the distribu-

tion of women's fertility types and the posterior probability that each woman belongs to each type from an auxiliary measurement system. The remaining steps condition on these objects. Subsection 4.2 uses across-market variation and the marriage market equilibrium conditions to recover the Pareto weights, adapting Gayle and Shephard (2019) to our setting. Subsection 4.3 then takes the Pareto weights as known and recovers the gender-specific preferences for children together with the remaining model primitives.

4.1 Unobserved Fertility Type

We introduce persistent unobserved heterogeneity only in women's preferences for children because the selection problem central to our analysis concerns how women with different fertility preferences sort into education. Specifically, each woman has a discrete *fertility type* $r \in \mathcal{R} = \{L, H\}$ that captures her intrinsic valuation of children. Men do not have an additional latent fertility type, although their preferences for children may differ from women's. Recovering the distribution of women's fertility types is necessary for identification of the remaining parameters. Because fertility type affects education, partnership, fertility, and labor supply choices, a specification that omits it may attribute differences in the composition of women making those choices to the structural parameters governing the choices themselves.

Persistent preference heterogeneity is also consequential for the counterfactual policy analysis. A model without such heterogeneity attributes all residual variation in fertility conditional on observables to transitory preference shocks. Such a model may place substantial probability mass near the fertility margin and therefore predict large responses to policies that reduce the cost of childbearing, including childcare subsidies, parental leave, and other work–family policies. Persistent heterogeneity can imply a different distribution of responses. High-type women may choose to have children even without the policy, while low-type women may place too little value on children for plausible cost reductions to change their fertility choices. Models with and without persistent heterogeneity can fit the same baseline fertility moments while producing sharply different policy responses. Baseline fit alone therefore cannot determine which specification is appropriate, motivating our direct use of auxiliary measurements to recover the type distribution.

Identification proceeds in two stages. The first stage recovers the share of each fertility type in market k and cohort c, $\pi_{r,k,c}$, and the posterior probability q_{ir} that woman i is of type r, conditional on her family background, auxiliary measurements, and observed outcomes. The second stage conditions on these type shares and posterior probabilities to recover the type-specific structural parameters, including the type-specific preferences for children, together with the remaining parameters of the discrete-choice and matching model. This subsection focuses on the first stage.

The first stage uses an auxiliary measurement system consisting of variables that do not enter the structural payoff equations but whose distributions vary with the latent fertility type. Together with flexible reduced-form equations for education, partnership, labor supply, and fertility, these measurements separate persistent type heterogeneity from transitory variation in observed choices. The measurements also give the latent types an economic interpretation (Carneiro et al., 2003). We use three measurements:

- **Age at first child.** Conditional on having a first child by age 38, a woman who places greater value on children has a stronger incentive to begin childbearing earlier because fecundity declines with age. We therefore allow the type-conditional distribution of age at first child to differ across fertility types, with high types giving birth earlier on average. Although the structural model does not explicitly model birth timing, this measurement allows differences in timing to inform the classification of women into fertility types. Birth timing therefore enters the analysis indirectly through the recovered type distribution.

- **Age at first relationship.** Conditional on entering a relationship by age 38, earlier partnership formation lengthens the period over which a woman can have children before fecundity declines. Women with stronger preferences for children may therefore enter relationships earlier. We allow the type-conditional distribution of age at first relationship to differ accordingly, with high types entering their first relationship earlier on average.

- **Number of siblings.** Number of siblings is informative about fertility type if preferences for children persist across generations. Previous literature has documented positive associations between family-of-origin size and fertility in the next generation (Murphy and Knudsen, 2002; Kolk, 2014; Keenan et al., 2025; Kwok, 2026). Because the number of siblings is determined before the woman's education, partnership, labor supply, and fertility choices, it provides a predetermined signal of her latent fertility type.

Each timing measurement contributes to the likelihood only when the corresponding event occurs by age 38. For a woman who does not enter a relationship or have a child by that age, non-occurrence represents an extensive-margin outcome rather than a late realization of the timing variable. Non-occurrence therefore enters the first stage through the partnership and number-of-children equations, while the corresponding timing measurement contributes no additional information. This treatment avoids counting the same extensive-margin outcome in multiple equations and preserves the conditional-independence structure imposed on the measurements. Appendix G provides details.

Given the likelihood structure and the conditions stated in Appendix G, the latent-class identification results of Allman et al. (2009) identify the population type shares and the type-conditional measurement distributions up to a relabeling of the types. The key restrictions are that, conditional on observed covariates and fertility type, the measurements are mutually independent and that their conditional distributions differ sufficiently across types. Intuitively, conditional on the covariates, fertility type must be the only source of dependence among the measurements. Their common dependence can then be attributed to the latent type rather than to measurement-specific noise. We label the type associated with earlier childbearing, earlier partnership formation, and a larger family of origin as the high fertility type. We then carry the identified type shares and posterior type probabilities into the identification of the remaining model parameters.

4.2 Pareto Weights

Following the equilibrium approach of Gayle and Shephard (2019) and Galichon et al. (2019), we identify Pareto weights from variation across marriage markets. Gayle and Shephard (2019) and Reynoso (2024) use across-market variation in the education distributions of women and men for this purpose. We cannot treat that variation as excluded here: education is chosen within our model and responds to wages, education costs, and expected marriage outcomes. We use variation in the sex ratio across local marriage markets instead (Bourguignon et al., 2009).

Let λ_k^{xy} denote the Pareto weight on the woman in a type-(x, y) couple in market k, and let $\mathbf{\Lambda} = \{\lambda_k^{xy}\}_{x\in\mathcal{E}\times\mathcal{R}, y\in\mathcal{E}, k\in\mathcal{K}}$ collect these weights. The sex ratio changes the relative supplies of women and men, which changes the weights required to clear the marriage market. Identification requires preferences, match-quality parameters, and shock distributions to be common across markets. Conditional on observed wages, policies, childcare costs, and divorce risk, the sex ratio must have no direct effect on household preferences or budgets. Under these restrictions, its effect on household choices operates through equilibrium Pareto weights.

Proposition 2. *Under the conditions of Proposition 1 and the restrictions across markets above, suppose the sex ratio varies locally while the observed market environment is held fixed and shifts the Pareto weight for each couple type. If matching and singlehood probabilities are positive, the vector of Pareto weights* $\mathbf{\Lambda}$ *is identified.*

Proof of Proposition 2. See subsection F.1. □

The intuition for identification is that relative scarcity changes the terms of marriage. When women are scarce relative to men, men must offer women more of the gains from marriage to attract a partner; this tends to raise wives' Pareto weights. When men are scarce, the pressure runs the other way. Under the exclusion restriction, sex ratios alter household choices through these equilibrium weights rather than directly through preferences or budgets.

4.3 Heterogeneous Preferences for Children and the Rest of the Model Parameters

In standard household models, spouses' preferences for children are not separately identified because the data reveal only the couple's joint fertility choice. Eckstein et al. (2019) state the difficulty plainly: "Of course, one could imagine individuals in a couple getting different utilities from a pregnancy decision, but we cannot infer such differences from the data, so we ignore them." Separating spouses' preferences therefore typically requires auxiliary information, such as surveys eliciting stated childbearing preferences (Doepke and Kindermann, 2019). Such surveys are often unavailable and are rarely linked to population-wide administrative records.

Our framework instead uses variation in Pareto weights across marriage markets. Once the Pareto weights are known, the household choice probabilities identify, for each couple type and child count, a Pareto-weighted combination of the spouses' valuations of children. Across markets, the same couple type faces different equilibrium Pareto weights. These weights generate different combinations of the same underlying valuations, allowing the female and male components to be separated when the weights are not all equal. The argument requires

spouses' preferences for children to be invariant across markets and at least two distinct Pareto weights for each couple type. A modified version of the semiparametric discrete-choice identification arguments in Matzkin (1992, 1993) then identifies the remaining primitives of the household-life stage.

Proposition 3. *Given the identified Pareto weights, and under additional assumptions stated in subsection F.2, all other model primitives in the household-life stage, including gender-specific preferences for children, are identified.*

Proof of Proposition 3. See subsection F.2. □

As an example of the identification argument, consider the work–family cost. Labor supply without children disciplines baseline work disutility. Conditional on wages, fertility type, and the identified Pareto weights, labor supply across number of children and education groups then identifies the additional cost of working while raising children. This cost matters for fertility and for the response to childcare and wage policies.

Finally, given the Pareto weights, fertility-type distribution, and household-life-stage primitives, observed education choices, couple-type match frequencies, and singlehood identify the remaining parameters of the education-choice and matching stages.

Proposition 4. *Given the identified Pareto weights, fertility type, and model primitives in the household-life stage, the deterministic component of the gender-education-specific cost $c_{g,e}$, match quality parameters ξ^{xy}, and the scale of the matching taste shock σ_ν in the education choice and matching stages are identified, up to a normalization of one education-cost level per gender; under an additional cross-situation variation condition (Assumption ID-8 in subsection F.3), the scale of the education-cost shock σ_{ec} is identified as well.*

Proof of Proposition 4. See subsection F.3. □

The matching and education parameters use different choice margins. Once household values are known, the difference between women's and men's adjusted match-versus-singlehood odds identifies the scale of matching tastes whenever their economic gains from that match differ. The level of those odds then identifies the match-quality term common to both spouses. Education shares identify cost differences at a given shock scale. Variation in the returns to education across fertility types, markets, and cohorts separates those costs from the common scale of education-cost shocks under Assumption ID-8.

5 Estimation

We first describe the parameters set outside estimation (Subsection 5.1). We then explain the two-step procedure: estimating the distribution of women's fertility types with the expectation-maximization (EM) algorithm (Subsection 5.2.1) and, conditional on those estimates, the remaining structural parameters by the generalized method of moments (GMM) (Subsection 5.2.2).

5.1 Parameters Set Outside Estimation

We define four education categories—primary, secondary, bachelor's degree, and master's/PhD—and accordingly set $|\mathcal{E}| = 4$. We assume two fertility types, $\mathcal{R} = \{L, H\}$, representing low and high fertility preferences, respectively.

Table 1 summarizes the parameters set outside the estimation. We adopt the square-root equivalence scale to set both economies of scale in consumption, ψ, and the fraction of household income allocated to children, $\kappa(n)$ (Eckstein et al., 2019). We set the part-time scaling parameters equal to the ratio of part-time to full-time hours, $\beta_g^{PT} = \alpha_g^{PT} = 0.75$, for both genders, making the disutility of labor and the work–family cost linear in hours. To reduce the number of parameters, we restrict men's linear work–family coefficient to be common across education, $\alpha_{1,m,e} = \alpha_{1,m}$. This restriction is consistent with the limited variation in men's labor supply across education groups and numbers of children. We allow the corresponding coefficient for women to vary by education because the work–family cost may differ across the jobs associated with different education levels. The exogenous probability of divorce is specified as a function of couple type and the number of children and is taken directly from the data; as Table A.3 shows, it decreases with the number of children.

The parameters governing the baseline policy environment, including taxation, childcare costs, and child support, are calibrated to the corresponding Danish institutions by following Jakobsen et al. (forthcoming). Appendix C describes their construction.

5.2 Parameters Estimated within the Model

In the second stage, we estimate the cohorts jointly, constraining education costs and the education taste-shock scale to be common across them. We use the cohort-specific estimates to assess which changes in model primitives account for the changes in marriage and fertility documented in Section 2. This exercise requires utility to be measured in the same units in both cohorts. First, monetary quantities are expressed in real terms, so the budget constraint is denominated in a common unit across cohorts. Second, the coefficient on the private-consumption component of utility is normalized to one throughout the model, fixing the scale of utility across cohorts. Together, these choices place the preference and cost parameters on a common utility scale, making them comparable across cohorts.

5.2.1 Estimation of Fertility Type

The first step estimates two objects used in the second stage: the share $\pi_{r,k,c}$ of women of fertility type r in local marriage market k and cohort c, and the probability q_{ir} that woman i belongs to type r. The coefficients of the first-stage likelihood are estimated jointly across cohorts, while the type shares vary by market and cohort. We follow the approach of Arcidiacono et al. (2025).

Because fertility type is unobserved, the first-stage likelihood takes a finite-mixture form in which each woman's contribution is a weighted sum of type-specific likelihoods, with weights given by the type shares $\pi_{r,k,c}$ in her market and cohort. We maximize the finite-mixture likelihood using the EM algorithm (Dempster et al., 1977; Arcidiacono and Jones, 2003), which

Table 1: Parameters Set Outside Estimation

		Value	Source
Budget constraint			
ψ	Economies of scale	$2 \times \frac{1}{\sqrt{2}} = \sqrt{2}$	Eckstein et al. (2019)
$\kappa(n)$	Fraction of consumption on children	$1 - \sqrt{2/(2+n)}$	Eckstein et al. (2019)
$w_{g,e,n}$	Wages	See Table A.1, Table A.2	Calculated with data
Labor and work–family (part-time)			
β_g^{PT}	Part-time/full-time labor-disutility ratio	0.75	Hours ratio (imposed)
α_g^{PT}	Part-time work–family-cost scaling	0.75	Hours ratio (imposed)
Partnership dissolution			
$d^{xy}(n)$	Dissolution rate	See Table A.3	Calculated with data

Notes: The table reports the parameters set outside the estimation, together with their values and sources. Wage schedules and dissolution probabilities are constructed separately for each cohort. The remaining calibrated parameters are held fixed across cohorts.

alternates between computing posterior type probabilities and updating the common likelihood parameters and market–cohort type shares.

The first-stage likelihood combines three measurements informative about fertility type—age at first relationship, age at first child, and number of siblings—with flexible reduced-form specifications for education, marriage, labor supply, and the number of children. The EM algorithm produces the market- and cohort-specific type shares $\pi_{r,k,c}$ and, for each woman i, the posterior type probability q_{ir}. The remaining first-stage parameters are treated as nuisance parameters and do not enter the second-stage estimation. Appendix G provides the implementation details.

In the second stage, we take the estimated $(q_{ir}, \pi_{r,k,c})$ as given. The type shares $\pi_{r,k,c}$ enter the equilibrium calculation for each market and cohort and therefore affect the model-implied moments, while the individual posterior probabilities q_{ir} are used as weights to construct fertility-type-specific data moments. The moment definitions below indicate where these weights enter.

5.2.2 Estimation of the Remaining Parameters

We estimate the structural parameters by GMM. The 32 Pareto weights in each market, one for each combination of male education, female education, and female fertility type ($4 \times 4 \times 2$), are equilibrium objects rather than free parameters. For each candidate parameter vector $\boldsymbol{\Theta}$, we solve the marriage market equilibrium to obtain the vector of Pareto weights $\boldsymbol{\Lambda}(\boldsymbol{\Theta})$, stacked across markets and couple types. We then construct the model-implied moments $\mathbf{m}_{model}(\boldsymbol{\Theta}, \boldsymbol{\Lambda}(\boldsymbol{\Theta}))$ and compare them with the corresponding data moments $\mathbf{m}_{data}$. The GMM estimator chooses the parameter vector that minimizes the weighted distance between the model and data moments:

$$\hat{\boldsymbol{\Theta}} = \arg\min_{\boldsymbol{\Theta}} \left[\mathbf{m}_{model}(\boldsymbol{\Theta}, \boldsymbol{\Lambda}(\boldsymbol{\Theta})) - \mathbf{m}_{data}\right]' \mathbf{W} \left[\mathbf{m}_{model}(\boldsymbol{\Theta}, \boldsymbol{\Lambda}(\boldsymbol{\Theta})) - \mathbf{m}_{data}\right] \tag{31}$$

where $\mathbf{W}$ is a positive definite weighting matrix. In light of the finite-sample concerns raised by Altonji and Segal (1996) regarding optimal weighting matrices, we use a diagonal weighting matrix.

To account for sampling variation in both the first-stage estimates and the data moments, we compute standard errors from a stratified bootstrap with $B = 250$ replications. Within each market–cohort cell, we draw women and men separately with replacement, keeping the number of observations of each gender fixed and thus preserving observed market sizes and sex ratios. In each replication b, we re-estimate the first stage to obtain $(q_{ir}^{(b)}, \pi_{r,k,c}^{(b)})$, recompute the data moments using $q_{ir}^{(b)}$ as weights, and re-estimate $\boldsymbol{\Theta}$ by GMM with $\pi_{r,k,c}^{(b)}$ entering the equilibrium calculation. The reported standard errors are the sample standard deviations of the corresponding bootstrap estimates.

Following the identification argument in Section 4, we target three sets of moments separately in each market.

Marriage market and education: The first set consists of the cell shares of matched couples by husband's education, wife's education, and wife's fertility type ($4 \times 4 \times 2 = 32$ moments), single women by education and fertility type ($4 \times 2 = 8$ moments), and single men by education (4 moments), for a total of 44 moments. Cells indexed by the wife's fertility type are constructed using the posterior probabilities q_{ir} as weights. Because the set includes both matched individuals and singles, summing across partner education and fertility type yields the education distribution for each gender.

Children: The second set consists of the shares with exactly one, two, and three or more children, by couple type and the wife's fertility type ($4 \times 4 \times 2 \times 3 = 96$ moments).

Labor supply and work–family interaction: The third set consists of the share working and the share working full-time conditional on working. For women, both outcomes are constructed by education, marital status, and the number of children, with divorced women counted as single ($2 \times 4 \times 2 \times 4 = 64$ moments). For married men, they are constructed by education and the number of children ($2 \times 4 \times 4 = 32$ moments); for single men, they are constructed by education ($2 \times 4 = 8$ moments). This set contains 104 moments in total.

6 Estimation Results

We use the cohort estimates to assess how changes in preferences, childbearing costs, and marriage market conditions account for the fertility decline. We first report the main estimates and assess the model's fit.

6.1 Parameter Estimates

Table 2 reports the parameter estimates for both cohorts. Utility is measured in units of log consumption, whose coefficient is normalized to one in both cohorts. The two cohorts are estimated jointly, with education-cost levels and the education taste-shock scale held common across them. We focus on the implied utility from children and the cost of combining work and childrearing.

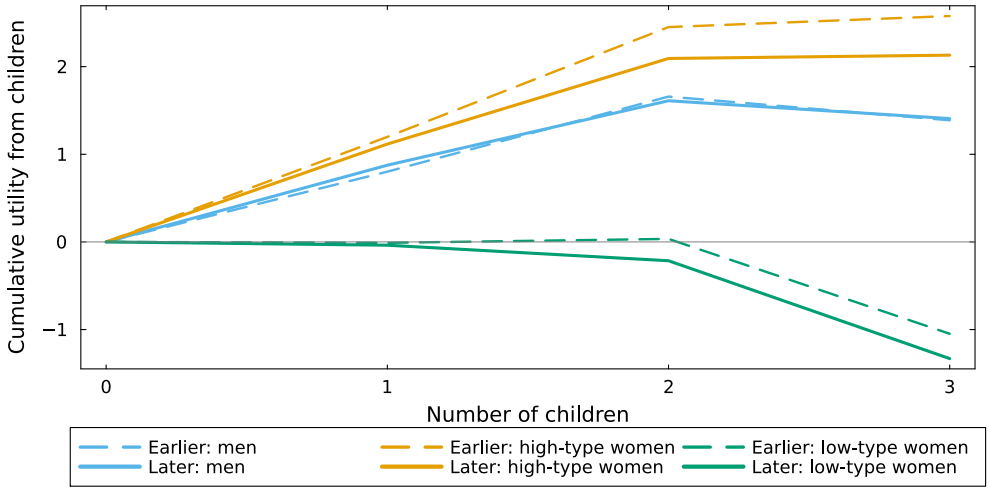
Table 2: Parameter estimates by cohort

		1967–69	1978–80
Utility from children (Figure 7)			
$\gamma_{1,m}$	Men, linear	2.547 (0.088)	2.144 (0.065)
$\gamma_{1,f,H}$	Women, high type, linear	2.944 (0.135)	2.385 (0.098)
$\gamma_{1,f,L}$	Women, low type, linear	1.735 (0.084)	1.231 (0.049)
γ_2	Curvature (common)	−0.564 (0.023)	−0.470 (0.018)
γ_0	Intercept (common)	−1.182 (0.057)	−0.799 (0.032)
Work–family cost at full-time work			
$\alpha_{0,f}$	Women, constant	0.150 (0.010)	0.099 (0.009)
$\alpha_{0,m}$	Men, constant	0.272 (0.017)	0.179 (0.017)
$\alpha_{1,f,P}$	Women, coeff. on n: primary	0.037 (0.025)	0.073 (0.006)
$\alpha_{1,f,S}$	secondary	0.004 (0.023)	0.000 (0.018)
$\alpha_{1,f,B}$	bachelor	0.019 (0.023)	0.001 (0.013)
$\alpha_{1,f,M}$	MA/PhD	0.047 (0.022)	0.014 (0.005)
$\alpha_{1,m}$	Men, coeff. on n (common across educ.)	0.006 (0.044)	0.000 (0.025)
$\alpha_{2,f}$	Women, coeff. on n^2	0.004 (0.006)	0.002 (0.005)
$\alpha_{2,m}$	Men, coeff. on n^2	0.002 (0.011)	0.014 (0.004)
Labor disutility (linear in hours; part-time $= 0.75\,\beta_g$)			
β_f	FT work disutility, women	−0.028 (0.028)	0.147 (0.008)
β_m	FT work disutility, men	0.356 (0.045)	0.398 (0.028)
Marriage stage			
ξ_0	Match-value intercept	−0.231 (0.053)	0.034 (0.041)
ξ_1	Mismatch, per level: husband more educated	−0.137 (0.013)	−0.208 (0.014)
ξ_2	Mismatch, per level: wife more educated	−0.129 (0.017)	−0.171 (0.015)
Shock scales			
σ_ν	Matching taste shock	0.169 (0.020)	0.232 (0.017)
σ_ε	Household choice shock	0.360 (0.011)	0.276 (0.007)
ρ_ε	Nest parameter	0.164 (0.006)	0.207 (0.008)
Common across cohorts			
$c_{f,P}$	Education cost, women: primary	0.087 (0.016)	
$c_{f,B}$	bachelor	0.119 (0.009)	
$c_{f,M}$	MA/PhD	0.394 (0.015)	
$c_{m,P}$	Education cost, men: primary	−0.005 (0.017)	
$c_{m,B}$	bachelor	0.437 (0.020)	
$c_{m,M}$	MA/PhD	0.572 (0.024)	
σ_{ec}	Education taste shock	0.100 (0.045)	
Set outside the second stage			
π_H	High-fertility-type share (first stage)	0.628 (0.005)	0.593 (0.008)

Notes: Estimates from the second-stage GMM described in Section 5, with the two cohorts estimated jointly. Standard errors in parentheses use 250 stratified bootstrap replications, with both estimation stages repeated in each replication. Utility is measured in units of log consumption (the consumption coefficient is normalized to one) in both cohorts. Education-cost levels are measured relative to secondary education and, together with the education taste-shock scale σ_{ec}, are common to both cohorts. The imposed parameters are listed in Table 1. π_H is estimated in the first stage and enters the second stage as an input. Location-averaged values are shown.

Preferences for children and work–family costs. Figure 7 reports estimated utility from children before work–family costs. At two children, this utility falls from 2.45 to 2.09 for high-fertility-type women and from 0.03 to −0.21 for low-fertility-type women. Weighting the types by their cohort-specific population shares gives a decline from 1.55 to 1.15 for women; this average reflects both changes in preferences and the lower high-type share. Men's utility at two children changes little, from 1.66 to 1.61. Men's utility rises slightly at one and three children.

Figure 7: Estimated utility from children by gender and fertility type



Notes: The figure plots estimated $\Gamma_{g,r}(n)$ before work–family costs. Dashed lines denote the earlier cohort and solid lines the later cohort. Women have high and low fertility types; the same preference function applies to all men.

The fixed component of the work–family cost also falls, from 0.15 to 0.10 for married women and from 0.27 to 0.18 for men. Per-child costs are small for most groups, but remain appreciable for primary-educated women in the later cohort (Table 2). These reductions partly offset the effect of women's declining utility from children. Net valuations also differ across female fertility types and education groups; Figure A.2 reports profiles averaged over fertility types by education.

Other parameters. The match-value intercept ξ_0 rises from -0.23 to 0.03. The decomposition below shows that changes in match value offset part of the fertility decline associated with women's preferences, mainly through partnership and matching. The estimated disutility of full-time work rises, particularly for women, and the dispersion of matching tastes increases. Education-cost levels are constrained to be common across cohorts; within this specification, changes in educational attainment reflect changes in the other estimated inputs and equilibrium returns. The first-stage estimate of the high-fertility-type share falls from 0.63 to 0.59. Wives' matching-weighted mean Pareto weight rises from 0.62 to 0.66 (Table A.4). For wives with at least secondary education, it rises in 23 of the 24 couple-type-by-fertility-type cells. This is an equilibrium outcome, but the present decomposition does not isolate which changes in primitives account for the rise in the weight.

6.2 Model Fit

We assess fit in education, partnership, fertility, and employment. Appendix A reports additional model fit.

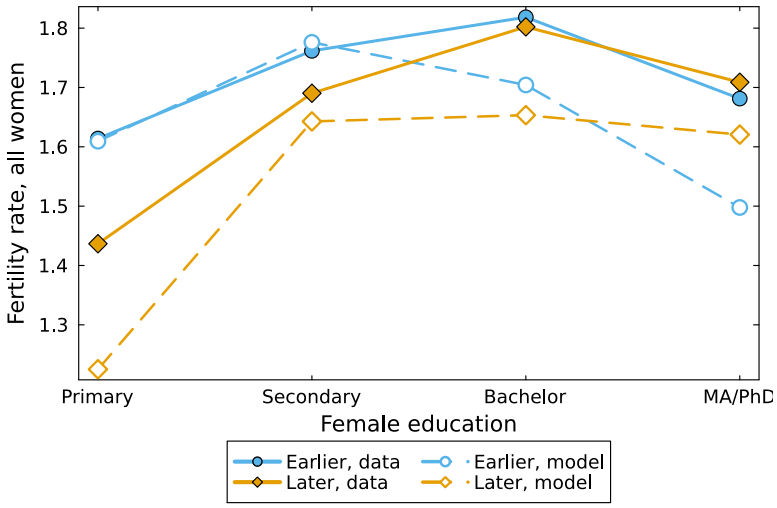
The model closely reproduces partnership and education: the aggregate partnership rate differs from the data by at most 0.6 percentage points in either cohort, and all education shares by gender and cohort are within 4 percentage points. Predicted fertility is 1.715 rather than

1.748 children per woman in the earlier cohort and 1.609 rather than 1.707 in the later cohort. The model therefore overstates the decline (0.105 against 0.040), primarily because it overstates the decline in fertility within couples.

Figure 8 shows that the model captures the fertility decline among primary-educated women and the near-stability among women with a bachelor's degree or more. The level differences are largest for master's/PhD women in the earlier cohort, whose fertility the model understates by 0.18 children, and for primary-educated women in the later cohort, for whom it overstates the decline by 0.21 children. The first difference partly reflects the model's overprediction of singlehood among master's/PhD women (Figure 9); the flatter within-couple education profile also contributes to the fertility gaps (Figure A.3).

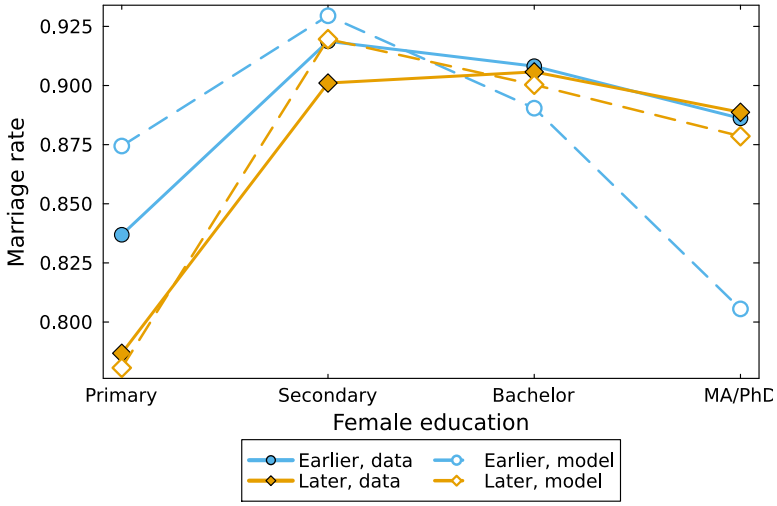
Figure 9 reports partnership rates by female education. The model reproduces the later-cohort profile within 2 percentage points: partnership is lowest among primary-educated women, and the largest across-cohort decline occurs at low education. In the earlier cohort it overstates singlehood among master's/PhD women (0.19 against 0.11) and understates it among primary-educated women (0.13 against 0.16). The model also reproduces the broad pattern of assortative matching by education in both cohorts (Table A.6).

Figure 8: Unconditional fertility rate by female education: data vs. model

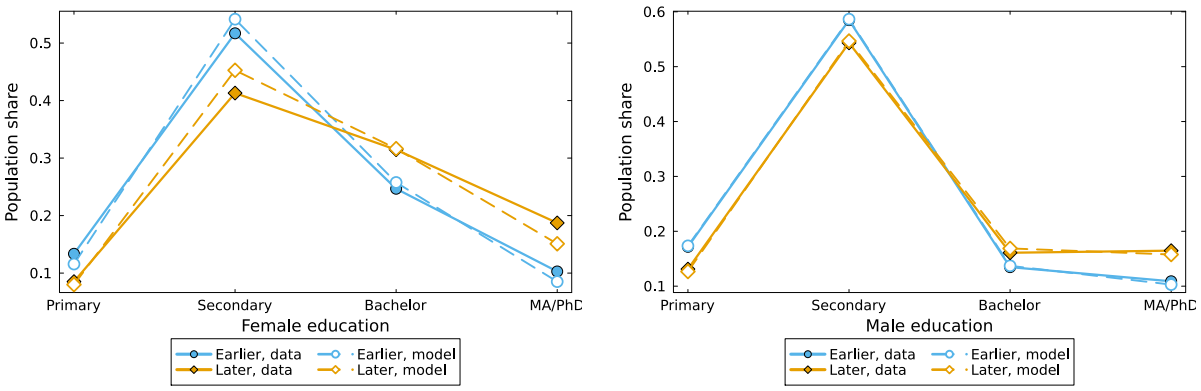


Notes: Data are from Danish administrative registers provided by Statistics Denmark; model rates use the estimated parameters. The earlier cohort comprises women born in 1967–69 and the later cohort women born in 1978–80. Fertility is the average number of children per woman accumulated between ages 19 and 38, including women without a qualifying partnership, and is reported by highest education completed by age 38. Filled markers with solid lines denote data and open markers with dashed lines model predictions; light blue denotes the earlier cohort and orange the later cohort.

Figure 9: Marriage rate by female education: data vs. model



Notes: Data are from Danish administrative registers provided by Statistics Denmark; model rates use the estimated parameters. The earlier cohort comprises women born in 1967–69 and the later cohort women born in 1978–80. The partnership rate is the share of women with a qualifying legal marriage or cohabiting partnership, as defined in Section 2.1, and is reported by highest education completed by age 38. Filled markers with solid lines denote data and open markers with dashed lines model predictions; light blue denotes the earlier cohort and orange the later cohort.



(a) Women

(b) Men

Figure 10: Education distribution by gender and cohort: data vs. model

Notes: Data are from Danish administrative registers provided by Statistics Denmark; model shares use the estimated parameters. The panels show population shares by highest education completed by age 38 for women and men born in 1967–69 (earlier cohort) and 1978–80 (later cohort). Filled markers with solid lines denote data and open markers with dashed lines model predictions; blue denotes the earlier cohort and orange the later cohort. Shares sum to one within gender and cohort.

Figure 10 shows that the model captures the larger increase in tertiary education among women than men. The largest gaps are for later-cohort women: the model slightly overstates the secondary share and understates the master's/PhD share.

For partnered women with secondary or higher education, most employment rates and full-time shares conditional on employment are within 5 percentage points of the data. Among partnered women, the model generally captures lower employment with more children and the earlier-cohort decline in full-time work, though it fits the later-cohort full-time profiles less closely (Table A.7 and Table A.8). Employment fit is weakest among primary-educated singles (Table A.7 and Table A.9).

As an external check, we compare the model's fertility response to a tax change with estimates from Danish tax reforms in Jakobsen et al. (forthcoming). For less-skilled couples, elasticities with respect to women's and men's net wages are -0.049 and 0.012 in the model, against -0.028 and 0.023 in their data (Table A.11). The signs and magnitudes are broadly similar, but the comparison is suggestive: their fertility measure covers four years, ours through age 38.

6.3 What Drives the Fertility Decline?

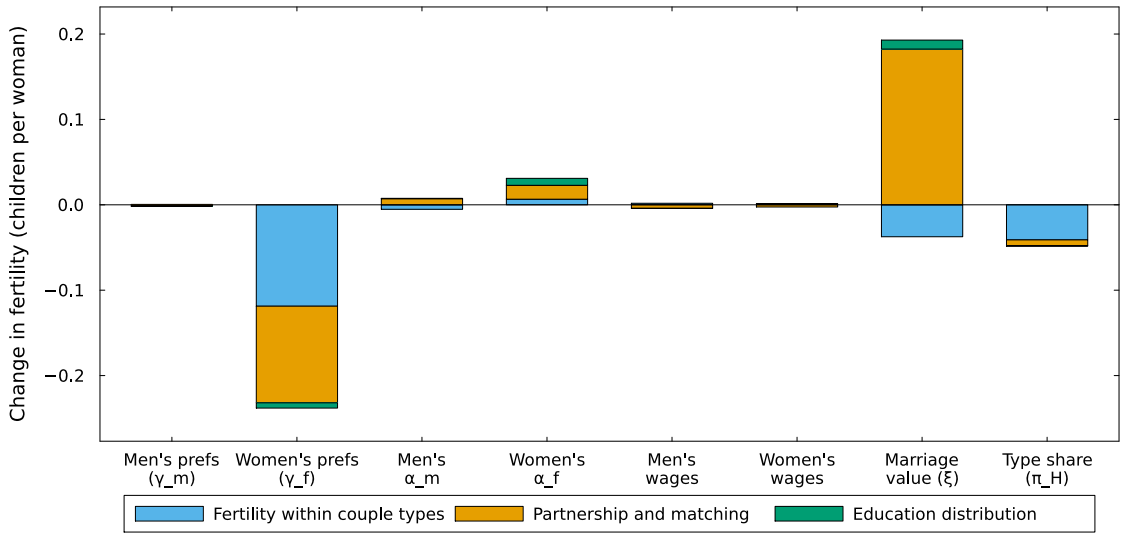
We decompose the model-predicted fertility decline by changes in nine blocks of parameters and inputs and by the margins through which those changes affect fertility. The blocks cover child preferences, work–family costs, wages, match value, the high-fertility-type share, and the remaining inputs. For each block, we average its effect over all orders of replacement and resolve the equilibrium at every step. We then divide that Shapley contribution among fertility within fixed couple types, partnership and matching, and women's education distribution. The partnership-and-matching margin includes both whether individuals form partnerships and who partners with whom. Figure 11 shows the eight named blocks; Table B.1 also reports the residual block and the exact values.

Changes in women's utility from children contribute -0.238 children per woman, equivalent to 227% of the model-predicted net decline of 0.105. That contribution exceeds the net decline because higher match value adds 0.156 children per woman, offsetting 149% of the decline, and lower work–family costs add 0.033, offsetting 31%. The lower high-fertility-type share contributes -0.048 (46%), while men's utility from children contributes -0.001 (1%) and the wage blocks together contribute -0.003 (3%).

Women's preference changes contribute -0.119 within couples and -0.113 through partnership and matching. Changes in match value add 0.182 through partnership and matching, more than offsetting the preference contribution on that margin. Summing all blocks, fertility within fixed couple types accounts for -0.101 (96%) of the model's -0.105 change; partnership and matching account for -0.003 (3%), and education for -0.001 (1%). The near-zero net contribution of the marriage market masks large, opposing responses to women's preferences and match value. Because the model overstates the observed within-couple decline, these allocations should not be read as measured contributions to the observed across-cohort change.

The model-implied education distributions of high- and low-fertility-type women are similar within each cohort (Table A.12). The estimates therefore provide little evidence that chang-

Figure 11: Shapley decomposition of the across-cohort fertility change by parameter block



Notes: Bars report parameter-block Shapley contributions to the across-cohort change in model-implied fertility per woman. Each bar is split among fertility within fixed couple types, partnership and matching, and women's education distribution. The figure omits the residual block, which is reported in Table B.1. All nine blocks sum to the model's decline of 0.105 children per woman. Section B.2 gives the construction.

ing selection of fertility types into education explains the relative stability of fertility among highly educated women.

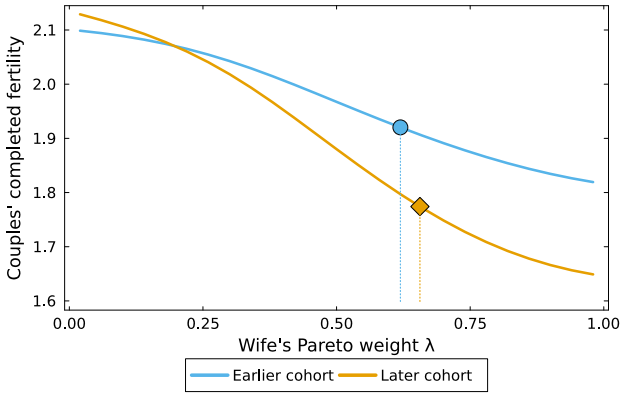
The parameter-block decomposition does not separately assign a contribution to changes in Pareto weights, which are determined in equilibrium. To assess their contribution to the across-cohort change, a second Shapley exercise varies household primitives, Pareto weights, and the masses of coupled and single individuals (subsection B.3). Changes in Pareto weights contribute $+0.007$ children per woman, compared with -0.064 for household primitives and -0.048 for couple composition. Wives' mean Pareto weight rises from 0.62 to 0.66 across the two cohorts (Table A.4).

The small across-cohort contribution of Pareto weights does not imply that fertility is insensitive to bargaining power. Figure 12 plots couples' average fertility as a function of the wife's Pareto weight, holding couple composition fixed. Raising the weight from zero to one lowers fertility by 0.28 children in the earlier cohort and 0.48 in the later cohort. The larger later-cohort response suggests that a policy-induced change in bargaining power can matter even though the estimated change in weights contributes little to the model-implied fertility decline. Section 7.4 examines this channel in the gender-wage counterfactual, which re-solves the marriage market and retains estimated divorce risk.

7 Policy Counterfactuals

We examine how education, marriage, and bargaining responses affect the fertility consequences of policy. In the later-cohort economy, we make childcare free, replace individual with joint taxation, raise women's wages, and subsidize tertiary education for each gender

Figure 12: Couples' fertility as a function of the wife's Pareto weight



Notes: Couples' completed fertility when every couple's Pareto weight is set to a common value λ (horizontal axis), with divorce risk set to zero. The exercise is partial equilibrium: each cohort's primitives and matched masses are held at its own baseline, each couple type's full choice problem is re-solved at the imposed λ, and the resulting fertility choices are aggregated with the baseline matched masses as weights. Markers indicate each cohort's estimated matching-weighted mean weight $\bar{\lambda}$ (Table A.4); their vertical coordinates are read from the common-λ schedule rather than the estimated equilibrium outcomes. Moving λ across its entire range changes completed fertility by 0.28 children in the earlier cohort and 0.48 in the later one.

separately. For each policy, we compare the full-equilibrium response with a benchmark that holds education and the marriage market fixed. We also report welfare separately for women and men to assess who benefits from each policy.

7.1 Policy Experiments and the Fixed-Couple Benchmark

For each policy, we calculate two responses. In the fixed-couple benchmark, each existing couple type re-optimizes labor supply and fertility at its original Pareto weight; education and the masses of couple and single types remain fixed. We then re-solve the equilibrium, allowing education, whether and whom individuals marry, and Pareto weights to change. Note that the fixed-couple benchmark holds the marriage market fixed and is not itself a new equilibrium.

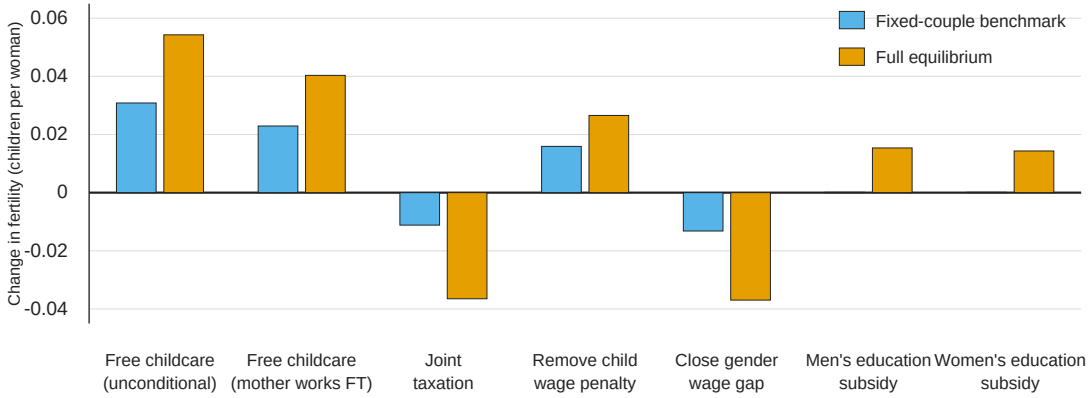
We evaluate private welfare separately for women and men, before they choose education and partners. Let V_g^0 and V_g^p denote mean ex ante inclusive utility for gender g in the baseline and policy equilibria. These values integrate over education-cost and matching shocks; for women, they also average over fertility types. With a unit coefficient on log consumption, the consumption-equivalent welfare change is

$$\mathrm{CEV}_g^p = 100 \left[\exp\left(V_g^p - V_g^0 \right) - 1 \right]. \tag{32}$$

This is the uniform percentage change in consumption that yields the same change in average ex ante utility for gender g as the policy. The main results exclude an offsetting fiscal adjustment; Appendix H reports supplementary financing exercises.

Figure 13 and Table 3 report the fertility effects and other outcomes of each policy. For childcare, taxation, and wages, the full-equilibrium fertility effect is larger in absolute value than the fixed-couple response. Education subsidies have no fixed-couple effect; their fertility

Figure 13: Policy effects on fertility: fixed-couple benchmark and full equilibrium



Notes: Changes in completed fertility per woman from the later-cohort baseline. The fixed-couple benchmark lets each baseline couple type re-optimize labor supply and fertility at its baseline Pareto weight and holds education and the population of every couple and single type fixed. The full-equilibrium calculation also allows education, partnership formation, partner composition, and Pareto weights to adjust. Education subsidies change no household budget, so their first bar is zero by construction. Table A.13 reports the plotted numbers.

effects arise through education and the marriage market. We now examine the mechanisms behind these differences for each policy.

7.2 Free Childcare and Partnership Formation

Childcare subsidies reduce the cost of children and can also increase the gains from forming a partnership. We set households' out-of-pocket childcare fees to zero, first unconditionally (Jakobsen et al., forthcoming) and then only when the mother works full-time. The unconditional exercise follows their free-childcare experiment; the full-time requirement introduces an additional incentive to work. Both policies leave existing child benefits unchanged.

The unconditional policy raises completed fertility by 0.031 children per woman in the fixed-couple benchmark and by 0.054 in full equilibrium. The benchmark captures existing couples' response to the lower cost of children. In full equilibrium, the prospect of raising children at lower cost also makes partnership more attractive. It raises the partnership rate by 1.5 percentage points. Using the same order-independent Shapley allocation as the descriptive and structural decompositions, Figure 14 assigns 41 percent of the fertility gain to partnership rates, 55 percent to fertility within couple types, and the remaining 4 percent to women's education and fertility-type composition and partner education. The work-contingent policy has a similar decomposition: it raises fertility by 0.023 in the fixed-couple benchmark and by 0.040 in full equilibrium, with 42 percent of the gain operating through partnership rates.

Low-fertility types account for a disproportionate share of the net increase in partnerships. Partnership rates are already about 95 percent among high-fertility types, compared with about 80 percent among low-fertility types. In the new equilibrium under unconditional free childcare, high-fertility types account for 46 percent of the net increase in partnerships, compared with 63 percent of baseline partnerships. Their smaller share in the expansion limits the fertility gain from higher partnership rates relative to a calculation that assigns the expan-

Table 3: Policy counterfactuals at the later-cohort baseline: full-equilibrium outcomes

				BA or above (pp)		Wife's	Welfare (CEV %)	
Policy	Fertility	Partnership (pp)	Female FT (pp)	Women	Men	Pareto weight	Women	Men
Family transfers								
Free childcare (unconditional)	+0.054	+1.5	−1.5	+0.8	+1.2	+0.007	+5.6	+3.0
Free childcare (mother works full-time)	+0.040	+1.1	+7.4	+1.5	+1.4	+0.014	+4.1	+2.2
Taxing the couple								
Joint taxation (income aggregation)	−0.037	−1.6	−4.9	−4.5	−5.7	0.000	−4.1	−2.7
Wage policies								
Remove motherhood wage penalty	+0.027	+0.6	+1.5	−0.3	−0.6	+0.006	+2.2	+1.2
Close gender wage gap	−0.037	−1.1	+3.1	+1.3	+0.4	+0.033	+12.0	−1.9
Education subsidies								
Men's education subsidy	+0.015	+0.5	+0.4	+11.4	+28.1	+0.007	+1.6	+9.6
Women's education subsidy	+0.014	+0.0	+1.2	+29.7	+15.4	+0.000	+19.1	+0.2
Later-cohort baseline (level)	1.609	89.6	81.2	46.7	30.5	0.656		

Notes: Each row reports changes from the later-cohort baseline after re-solving the full equilibrium, so education, partnership formation, partner composition, and Pareto weights all respond. "Fertility" is completed fertility per woman; "Partnership" is the economy-wide partnership rate; "Female FT" is the full-time share among working women; "BA or above" is the share with a bachelor's degree or above; and "Wife's Pareto wt." is the matching-weighted mean λ^{xy}. Welfare is the gender-specific consumption-equivalent change defined in (32); no offsetting fiscal adjustment is imposed. Appendix H reports supplementary revenue-neutral exercises. The last row reports the later-cohort baseline levels of each outcome, with the partnership rate, the female full-time share, and the education shares in percent so that they are comparable with the percentage-point changes above them.

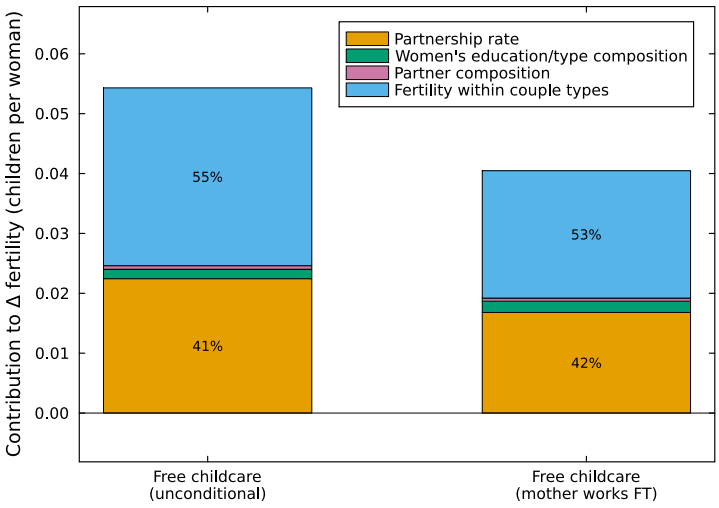
sion the baseline fertility-type composition.

Conditioning support on full-time work increases the financial return to full-time rather than part-time employment or nonemployment. The full-time share among working women rises by 7.4 percentage points under the conditional policy, compared with a decline of 1.5 points under unconditional provision. Unconditional provision lowers the cost of children without requiring employment, allowing some mothers to reduce work as family resources increase. Estimated work–family costs are lower for most women in the later cohort than in the earlier cohort, making full-time work less costly for mothers with children. Consistent with the lower cost of combining work and children, the full-time requirement increases work intensity while retaining about three-quarters of the unconditional fertility gain. The choice between the policies therefore depends on the weight placed on raising fertility relative to encouraging mothers' full-time work.

7.3 Joint Taxation

The tax treatment of couples affects both the return to work within marriage and the financial gains from marriage itself. We replace Denmark's individual income taxation with joint taxation under income aggregation, a form of joint taxation also studied by Gayle and Shephard (2019). The spouses' incomes are combined and the existing individual schedule is applied to the sum, so the secondary earner's first krone is taxed at the primary earner's marginal rate.

Figure 14: Shapley decomposition of childcare effects on fertility



Notes: Each bar decomposes the change in completed fertility per woman using four factors. Writing fertility as $\sum_{x,r} s_{xr} p_{xr} \sum_y \pi_{y|xr} n_{xyr}$, s_{xr} is the distribution of women across education x and fertility type r, p_{xr} is their partnership rate, $\pi_{y|xr}$ is the distribution of partner education y, and n_{xyr} is expected fertility within a couple type. The Shapley value averages each factor's marginal contribution across all possible orderings, so the four components sum to the total policy effect.

Joint taxation reduces the return to earnings within dual-earner couples and lowers the value of these partnerships. Under a progressive schedule, combining incomes raises the tax cost of a second earner relative to individual taxation and can make remaining single more attractive. It therefore changes the expected return to education for both genders, although women are the secondary earners in 82 percent of baseline couples. Tertiary attainment falls by 4.5 percentage points among women and 5.7 points among men. The partnership rate falls by 1.6 percentage points, the husband–wife education correlation falls by 0.03, and the full-time share among working women falls by 4.9 points. The lower partnership rate and weaker educational sorting are consistent with Gayle and Shephard (2019), who obtain both responses under income aggregation relative to their unrestricted optimal tax system.

These marriage and education responses amplify the fertility decline. The fixed-couple benchmark predicts a decline of 0.011 children per woman, whereas the full equilibrium produces a decline of 0.037, more than three times as large. The benchmark allows labor supply to adjust but omits the loss of partnerships and the shift toward lower education and earnings. Among baseline couple types, fertility falls most when both spouses are highly educated (Table A.15). Higher primary earnings expose more of the secondary earner's income to the upper tax bracket, and a highly educated secondary earner has more earnings at stake. The resulting loss of disposable income discourages childbearing, although the lower return to work also reduces its opportunity cost. At the estimated parameters, the net fertility response ranges from essentially zero for primary-educated couples to -0.028 children for couples in which both spouses have master's or doctoral degrees. Joint taxation reduces the value of dual-earner marriage, turning a modest fertility decline within existing couples into a much larger aggregate decline.

7.4 Wage Policies and Bargaining

Higher wages for women raise household resources but can also change the opportunity cost of children, the gains from marriage, and bargaining power within couples. We distinguish the level gap between childless men's and women's wages from the wage penalty associated with motherhood (Kleven et al., 2019). To close the level gap, we raise women's wages at every number of children by the ratio of men's to women's childless wages, preserving the proportional motherhood penalty. To remove the motherhood penalty, we set women's wages with children equal to childless women's wages at the same education level. The wage-level exercise holds men's wages fixed, isolating the household and marriage market responses to higher female wages.

Closing the level gap lowers completed fertility by 0.037 children per woman, despite the increase in women's earnings. The income gain can support more children, but higher wages also increase earnings forgone when mothers reduce work and enlarge the absolute wage penalty associated with motherhood. Moreover, higher wages improve women's opportunities in singlehood, reducing their dependence on marriage for income. The partnership rate falls by 1.1 percentage points, and wives' mean Pareto weight rises by 0.033, or about 5 percent of its baseline value of 0.656 (Table A.4). This increase gives greater weight to wives' choices; at the later-cohort estimates, wives' average valuation of children is below husbands'.

At baseline Pareto weights, the wage increase lowers fertility by 0.013 children per woman after couples adjust labor supply and fertility to their new budgets. Imposing the Pareto weights from the new equilibrium under the policy on these same baseline couples increases the decline to 0.020. The additional 0.007 decline is the bargaining contribution under this ordering of adjustments. Allowing education and marriage outcomes to adjust produces the full-equilibrium decline of 0.037.

Removing the motherhood penalty instead raises fertility by 0.027 children per woman and the partnership rate by 0.6 percentage points. Unlike the wage-level increase, this policy raises earnings when women have children and therefore improves the financial return to combining motherhood and work relative to remaining childless. It also increases the value of a partnership in which children are anticipated. Both wage policies raise the full-time share among working women, by 3.1 percentage points when the level gap is closed and 1.5 points when the motherhood penalty is removed. Their opposite fertility effects show why the source of a reduction in the gender wage gap matters for family outcomes.

7.5 Education Subsidies and the Marriage Market

Education subsidies can affect fertility through earnings and the marriage market, even though they do not directly change existing couples' budgets. We consider reductions in the costs of bachelor's and master's/PhD education for men and women separately. Women are more educated than men in the baseline, and the estimated costs of tertiary education are higher for men. The first policy sets men's costs equal to women's, reductions of 0.318 and 0.178 in log-consumption units. For individuals choosing the subsidized degrees, these reductions are equivalent to consumption increases of approximately 37 and 20 percent. The second policy

applies the same cost reductions to women. Table 4 reports the full-equilibrium results.

Both subsidies raise completed fertility by about 0.015 children per woman, but through different margins. The men's subsidy raises the partnership rate by 0.47 percentage points and average fertility among couples by 0.008 children; increased partnership accounts for about half of the aggregate fertility gain. The women's subsidy raises the partnership rate by only 0.04 percentage points and average fertility among couples by 0.015 children. Both policies also induce education among those who receive no subsidy: women's tertiary attainment rises by 11.4 percentage points when men are subsidized, and men's rises by 15.4 points when women are subsidized. More-educated potential partners raise the return to one's own education. Given preferences for similarly educated partners, subsidizing men also closes the initial gender education gap and improves partnership prospects for highly educated women. Subsidizing women widens that gap, and the partnership rate changes little. The nearly unchanged share of same-education couples under the men's subsidy masks a shift in partner composition: the share in which the wife is more educated falls by 7.4 percentage points, while the share in which the husband is more educated rises by 7.6 points.

Average fertility among couples rises under both policies as matches shift toward education pairings with higher fertility. The increase is smaller under the men's subsidy for two reasons. Women drawn into partnerships have lower expected fertility than women already partnered. In addition, the mean Pareto weight among wives with at least a bachelor's degree rises by 0.012 under the men's subsidy; at the later-cohort estimates, greater bargaining power for these wives tends to reduce fertility. The men's subsidy therefore gains more fertility through partnership formation, while the women's subsidy gains more through average fertility among couples.

The welfare effects also depend on which gender is subsidized. Subsidizing men raises women's welfare by 1.6 percent despite no direct reduction in their education costs, while men's welfare rises by 9.6 percent. Subsidizing women raises women's welfare by 19.1 percent but men's by only 0.2 percent. These comparisons involve large predicted changes in attainment: tertiary attainment rises by 28.1 percentage points among subsidized men and 29.7 points among subsidized women. The welfare results therefore rely on the model's predictions for education changes larger than those observed across cohorts.

7.6 Welfare

An increase in fertility need not imply a welfare gain, and a policy can improve welfare while reducing fertility. Individuals value consumption and leisure as well as children, and policies can affect women's and men's welfare differently. Policy evaluation therefore requires considering fertility alongside the welfare of both genders. Figure 15 reports gender-specific consumption-equivalent changes, and Table H.2 reports them by education.

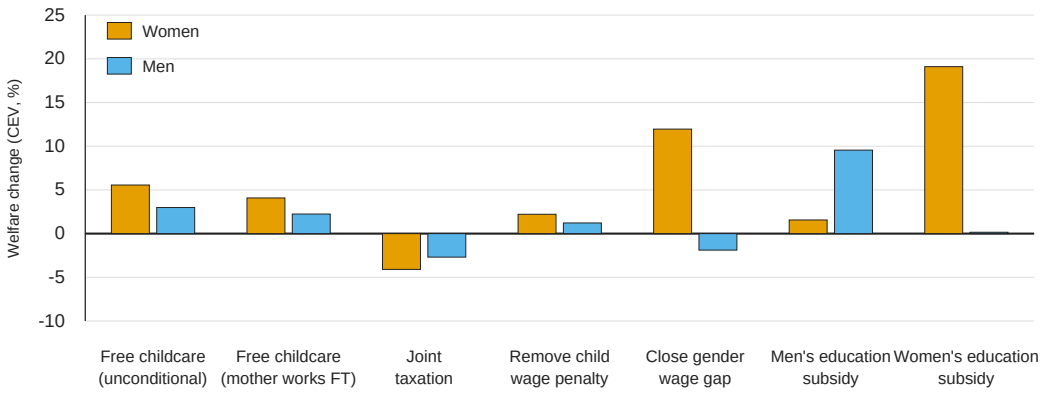
Free childcare and removal of the motherhood wage penalty raise both fertility and welfare for women and men by reducing the costs associated with children. Unconditional childcare raises women's welfare by 5.6 percent and men's by 3.0 percent; the full-time-contingent policy yields smaller gains of 4.1 and 2.2 percent because it restricts eligibility. Joint taxation instead lowers welfare for both genders, as higher taxes on dual-earner couples reduce earnings in-

Table 4: Education subsidies by gender: full-equilibrium outcomes

	Subsidized gender	
	Men	Women
Δ Fertility, all women	+0.015	+0.014
Δ Fertility, couples	+0.008	+0.015
Δ Partnership rate (pp)	+0.47	+0.04
Δ Same-education couples (pp)	−0.1	−6.8
Δ Wife-more-educated couples (pp)	−7.4	+6.6
Δ Husband-more-educated couples (pp)	+7.6	+0.2
Δ Wife's Pareto weight, all couples	+0.007	+0.000
wives with BA or above	+0.012	−0.011
Δ BA or above, women (pp)	+11.4	+29.7
Δ BA or above, men (pp)	+28.1	+15.4
Welfare, women (CEV %)	+1.6	+19.1
Welfare, men (CEV %)	+9.6	+0.2

Notes: Changes from the later-cohort baseline after re-solving the full equilibrium. "Men" sets men's bachelor's and master's/PhD education costs equal to women's estimated costs, reductions of 0.318 and 0.178 in log-consumption units; "Women" applies the same reductions to women's costs. The same-education, wife-more-educated, and husband-more-educated shares summarize partner composition. "Wife's Pareto weight" is the matching-weighted mean λ^{xy}, over all couples and over couples whose wife holds a bachelor's degree or above. The welfare rows report the gender-specific consumption-equivalent changes defined in (32); no offsetting fiscal adjustment is imposed. Appendix H reports supplementary revenue-neutral exercises.

Figure 15: Welfare effects by gender



Notes: Gender-specific welfare changes in consumption-equivalent percent from the later-cohort baseline, with the full equilibrium re-solved. These calculations do not impose an offsetting fiscal adjustment.

centives and the gains from partnership. Closing the gender wage gap illustrates a different trade-off: fertility falls, but women's welfare rises by 12.0 percent, while men's falls by 1.9 percent. Women benefit from higher earnings both within and outside marriage and gain bargaining power within couples. Men can lose despite unchanged wages and higher household income because they receive a smaller share of the gains within marriage and face a lower partnership rate.

Education subsidies generate welfare effects beyond their direct recipients. Subsidizing men raises women's welfare by 1.6 percent without directly lowering their education costs: the two education distributions become more similar, improving women's marriage prospects. Subsidizing women widens the initial education gap and produces little welfare gain for men,

despite their increase in attainment. The welfare benefit to the unsubsidized gender thus depends on how the policy changes the distribution of potential spouses, not only on the earnings gains of its recipients.

These welfare calculations exclude an offsetting fiscal adjustment. Appendix H reports supplementary exercises that restore baseline net government revenue with a common per-adult lump-sum tax. Under this financing rule, childcare contingent on mothers' full-time work still raises fertility and both genders' welfare, while increasing the full-time share among working women by 7.9 percentage points. Restricting eligibility lowers its direct fiscal cost, and the additional full-time work generates tax revenue that helps finance the policy. The combined responses of taxes and existing transfers recover 54 percent of its direct cost, compared with 21 percent for unconditional provision.

8 Conclusion

A descriptive decomposition of two Danish cohorts assigns nearly as much of the observed fertility decline to reduced partnership formation as to lower fertility within couples, while educational expansion offsets part of the decline. To understand the forces behind these changes, we estimate a model in which education, marriage, and fertility are jointly determined. Variation in local marriage market conditions, together with couples' joint choices, allows us to identify women's and men's preferences for children separately.

The model attributes its predicted cohort decline primarily to a fall in women's preferences for children; men's preferences change little. Lower work–family costs and a higher value of marriage partly offset this shift. Changes in women's preferences reduce fertility both within couples and through partnership and matching, but the higher value of marriage more than offsets the latter contribution. The net contribution of partnership and matching to the model-implied fertility decline is therefore small, despite the large partnership contribution in the descriptive accounting. Changes in Pareto weights also contribute little to that decline, although fertility is more sensitive to bargaining power under the later-cohort estimates.

These small net contributions to the cohort change do not make marriage or bargaining unimportant for policy. Free childcare induces new partnerships, but women with lower fertility preferences account for a disproportionate share of the increase. Joint taxation discourages education and marriage, making its fertility effect more than three times its fixed-couple effect. Education subsidies raise fertility through earnings and marriage market responses; subsidizing men also raises women's welfare by narrowing the gender education gap. Closing the gender wage gap lowers fertility partly by shifting bargaining power toward wives, whereas removing the motherhood wage penalty raises it. In the wage-level experiment, women's private welfare rises even as fertility falls.

References

Adda, Jérôme, Christian Dustmann, and Katrien Stevens, "The career costs of children," *Journal of Political Economy*, 2017, *125* (2), 293–337.

Allman, Elizabeth S, Catherine Matias, and John A Rhodes, "Identifiability of parameters in latent structure models with many observed variables," 2009.

Almar, Frederik, Benjamin Friedrich, Ana Reynoso, Bastian Schulz, and Rune M Vejlin, "Families' Career Investments and Firms' Promotion Decisions," Technical Report, National Bureau of Economic Research 2025.

Altonji, Joseph G and Lewis M Segal, "Small-sample bias in GMM estimation of covariance structures," *Journal of Business & Economic Statistics*, 1996, *14* (3), 353–366.

Arcidiacono, Peter and John Bailey Jones, "Finite mixture distributions, sequential likelihood and the EM algorithm," *Econometrica*, 2003, *71* (3), 933–946.

_ **and Robert A Miller**, "Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity," *Econometrica*, 2011, *79* (6), 1823–1867.

_ **, Esteban Aucejo, Arnaud Maurel, and Tyler Ransom**, "College attrition and the dynamics of information revelation," *Journal of Political Economy*, 2025, *133* (1), 53–110.

Bailey, Martha J, "The economics of childbearing: Trends, progress, and challenges," *Annual Review of Economics*, 2025, *17* (1), 695–720.

Baudin, Thomas, David De La Croix, and Paula E Gobbi, "Fertility and childlessness in the United States," *American Economic Review*, 2015, *105* (6), 1852–1882.

BEF, Statistics Denmark, "BEF - Befolkningen, Statistics Denmark: Population and Education, Social statistics.," 1980–2018. Data accessed through the ECONAU project database.

Bick, Alexander, "The quantitative role of child care for female labor force participation and fertility," *Journal of the European Economic Association*, 2016, *14* (3), 639–668.

Bisin, Alberto and Giulia Tura, "Marriage, fertility, and cultural integration in Italy," *American Economic Review*, 2026, *116* (9), 3423–3463.

Bloom, David E, Michael Kuhn, and Klaus Prettner, "Fertility in high-income countries: Trends, patterns, determinants, and consequences," *Annual Review of Economics*, 2024, *16*.

Bourguignon, François, Martin Browning, and Pierre-André Chiappori, "Efficient intra-household allocations and distribution factors: Implications and identification," *The Review of Economic Studies*, 2009, *76* (2), 503–528.

Carneiro, Pedro, Karsten T Hansen, and James J Heckman, "Estimating distributions of treatment effects with an application to the returns to schooling and measurement of the effects of uncertainty on college choice," *International Economic Review*, 2003, *44* (2), 361–422.

Caucutt, Elizabeth M, Nezih Guner, and John Knowles, "Why do women wait? Matching, wage inequality, and the incentives for fertility delay," *Review of Economic Dynamics*, 2002, *5* (4), 815–855.

Chiappori, Pierre-André and Bernard Salanié, "Mating markets," in "Handbook of the Economics of the Family," Vol. 1, Elsevier, 2023, pp. 49–109.

_ , _ , and Yoram Weiss, "Partner choice, investment in children, and the marital college premium," *American Economic Review*, 2017, *107* (8), 2109–2167.

_ , Monica Costa Dias, and Costas Meghir, "The marriage market, labor supply, and education choice," *Journal of Political Economy*, 2018, *126* (S1), S26–S72.

_ , Murat Iyigun, and Yoram Weiss, "Investment in schooling and the marriage market," *American Economic Review*, 2009, *99* (5), 1689–1713.

Choo, Eugene and Aloysius Siow, "Who marries whom and why," *Journal of Political Economy*, 2006, *114* (1), 175–201.

Dempster, Arthur P, Nan M Laird, and Donald B Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *Journal of the royal statistical society: series B (methodological)*, 1977, *39* (1), 1–22.

Doepke, Matthias and Fabian Kindermann, "Bargaining over babies: Theory, evidence, and policy implications," *American Economic Review*, 2019, *109* (9), 3264–3306.

_ , Anne Hannusch, Fabian Kindermann, and Michèle Tertilt, "The economics of fertility: A new era," in "Handbook of the Economics of the Family," Vol. 1, Elsevier, 2023, pp. 151–254.

Eckstein, Zvi, Michael Keane, and Osnat Lifshitz, "Career and family decisions: Cohorts born 1935–1975," *Econometrica*, 2019, *87* (1), 217–253.

Erosa, Andrés, Luisa Fuster, and Diego Restuccia, "A general equilibrium analysis of parental leave policies," *Review of Economic Dynamics*, 2010, *13* (4), 742–758.

Francesconi, Marco, "A joint dynamic model of fertility and work of married women," *Journal of Labor Economics*, 2002, *20* (2), 336–380.

Galichon, Alfred and Simon Weber, "Matching under imperfectly transferable utility," in "Handbook of the Economics of Matching," Vol. 1, Elsevier, 2024, pp. 79–122.

_ , Scott Duke Kominers, and Simon Weber, "Costly concessions: An empirical framework for matching with imperfectly transferable utility," *Journal of Political Economy*, 2019, *127* (6), 2875–2925.

Gayle, George-Levi and Andrew Shephard, "Optimal taxation, marriage, home production, and family labor supply," *Econometrica*, 2019, *87* (1), 291–326.

_ , Andrés Hincapié, and Robert A Miller, "Life-Cycle Fertility, Human Capital, and Family Policies: A Discrete-Continuous Choice Framework," *Working Paper*, 2024.

Greenwood, Jeremy, Nezih Guner, Georgi Kocharkov, and Cezar Santos, "Technology and the changing family: A unified model of marriage, divorce, educational attainment, and married female labor-force participation," *American Economic Journal: Macroeconomics*, 2016, *8* (1), 1–41.

Guner, Nezih, Ezgi Kaya, and Virginia Sánchez-Marcos, "Labor market institutions and fertility," *International Economic Review*, 2024, *65* (3), 1551–1587.

Gutierrez, Federico H, "A simple solution to the problem of independence of irrelevant alternatives in Choo and Siow marriage market model," *Economics Letters*, 2020, *186*, 108785.

Ho, Christine and Yutao Wang, "Social institutions and low birth rates," 2024.

IND, Statistics Denmark, "IND - Indkomst, Statistics Denmark: Labour and Income, Social statistics.," 1980–2018. Data accessed through the ECONAU project database.

Jakobsen, Katrine Marie, Thomas Høgholm Jørgensen, and Hamish Low, "Fertility and family labor supply," *American Economic Review*, forthcoming.

Kearney, Melissa S and Phillip B Levine, "Why Is Fertility So Low in High-Income Countries?," *Journal of Economic Literature*, 2026, *64* (3), 907–949.

Keenan, Katherine, Júlia Mikolai, Rebecca King, and Hill Kulu, "Intergenerational transmission of fertility in Great Britain: A parity-specific investigation using the 1970 British Cohort Study," *Population Studies*, 2025, *79* (3), 445–462.

Kim, Seongeun, Michele Tertilt, and Minchul Yum, "Status externalities in education and low birth rates in Korea," *American Economic Review*, 2024, *114* (6), 1576–1611.

Kleven, Henrik, Camille Landais, and Jakob Egholt Søgaard, "Children and gender inequality: Evidence from Denmark," *American Economic Journal: Applied Economics*, 2019, *11* (4), 181–209.

Kolk, Martin, "Multigenerational transmission of family size in contemporary Sweden," *Population studies*, 2014, *68* (1), 111–129.

Kwok, Tsz Chun, "Family Fertility Norms," 2026.

Lund, Christian Giødesen and Rune Majlund Vejlin, "Documenting and improving the hourly wage measure in the Danish IDA Database," 2016.

Matzkin, Rosa L, "Nonparametric and distribution-free estimation of the binary threshold crossing and the binary choice models," *Econometrica*, 1992, pp. 239–270.

_ , "Nonparametric identification and estimation of polychotomous choice models," *Journal of Econometrics*, 1993, *58* (1-2), 137–168.

Murphy, Michael and Lisbeth B Knudsen, "The intergenerational transmission of fertility in contemporary Denmark: The effects of number of siblings (full and half), birth order, and whether male or female," *Population studies*, 2002, *56* (3), 235–248.

Nöldeke, Georg and Larry Samuelson, "Investment and competitive matching," *Econometrica*, 2015, *83* (3), 835–896.

OECD, "Fertility rates," https://www.oecd.org/en/data/indicators/fertility-rates.html 2026. OECD Data Indicator. Accessed July 5, 2026.

RAS, Statistics Denmark, "RAS - Registerbaseret arbejdsstyrkestatistik, Statistics Denmark: Labour and Income, Social Statistics.," 1980–2018. Data accessed through the ECONAU project database.

Reynoso, Ana, "The impact of divorce laws on the equilibrium in the marriage market," *Journal of Political Economy*, 2024, *132* (12), 4155–4204.

Sommer, Kamila, "Fertility choice in a life cycle model with idiosyncratic uninsurable earnings risk," *Journal of Monetary Economics*, 2016, *83*, 27–38.

UDDA, Statistics Denmark, "UDDA - Højest fuldført uddannelse, Statistics Denmark: Population and Education, Social statistics.," 1980–2018. Data accessed through the ECONAU project database.

Wang, Hanna, "Fertility and family leave policies in Germany: Optimal policy design in a dynamic framework," *Review of Economic Studies*, forthcoming.

Yamaguchi, Shintaro, "Effects of parental leave policies on female career and fertility choices," *Quantitative Economics*, 2019, *10* (3), 1195–1232.

Appendices

A Additional Tables and Figures

Table A.1: Mean Log Wages by Education: Men

	Cohort 1967–69				Cohort 1978–80			
	Primary	Secondary	Bachelor	Master/PhD	Primary	Secondary	Bachelor	Master/PhD
for all n	5.801	5.907	6.050	6.191	5.854	5.946	6.025	6.181

Notes: The table reports mean log wages for men by education measured over ages 30–38. Wages are calculated from the Danish income register (IND, Statistics Denmark, 1980–2018) and enter the model as inputs.

Table A.2: Mean Log Wages by Education and Number of Children: Women

	Cohort 1967–69				Cohort 1978–80			
	Primary	Secondary	Bachelor	Master/PhD	Primary	Secondary	Bachelor	Master/PhD
$n = 0$	5.623	5.723	5.792	5.995	5.668	5.765	5.863	6.018
$n = 1$	5.579	5.669	5.726	5.975	5.668	5.723	5.824	6.002
$n = 2$	5.579	5.645	5.721	5.964	5.669	5.726	5.823	6.022
$n = 3$	5.539	5.600	5.693	5.949	5.613	5.680	5.804	6.006

Notes: The table reports mean log wages for women by education and number of children measured over ages 30–38. Wages are calculated from the Danish income register (IND, Statistics Denmark, 1980–2018) and enter the model as inputs.

Table A.3: Divorce probability by couple type and number of children

	Cohort 1967–69				Cohort 1978–80			
Couple type (Male, Female)	$n=0$	$n=1$	$n=2$	$n=3$	$n=0$	$n=1$	$n=2$	$n=3$
Primary, Primary	0.61	0.52	0.38	0.37	0.60	0.55	0.40	0.35
Primary, Secondary	0.60	0.49	0.33	0.30	0.68	0.52	0.35	0.29
Primary, Bachelor	0.69	0.43	0.30	0.21	0.62	0.43	0.32	0.23
Primary, Master/PhD	0.71	0.42	0.24	0.03	0.53	0.33	0.32	0.19
Secondary, Primary	0.53	0.42	0.28	0.27	0.54	0.44	0.29	0.26
Secondary, Secondary	0.64	0.40	0.27	0.23	0.67	0.37	0.23	0.19
Secondary, Bachelor	0.54	0.27	0.16	0.12	0.57	0.26	0.17	0.12
Secondary, Master/PhD	0.60	0.21	0.14	0.06	0.53	0.30	0.11	0.06
Bachelor, Primary	0.57	0.47	0.27	0.23	0.51	0.43	0.24	0.19
Bachelor, Secondary	0.66	0.38	0.23	0.19	0.66	0.33	0.18	0.13
Bachelor, Bachelor	0.58	0.29	0.16	0.11	0.58	0.26	0.14	0.10
Bachelor, Master/PhD	0.57	0.26	0.12	0.06	0.59	0.22	0.10	0.05
Master/PhD, Primary	0.60	0.40	0.26	0.25	0.66	0.37	0.21	0.14
Master/PhD, Secondary	0.60	0.30	0.20	0.14	0.62	0.24	0.13	0.07
Master/PhD, Bachelor	0.60	0.23	0.13	0.07	0.60	0.22	0.09	0.05
Master/PhD, Master/PhD	0.62	0.22	0.10	0.09	0.66	0.18	0.07	0.04

Notes: Data are from Danish administrative registers provided by Statistics Denmark. The table reports the probability that a qualifying partnership dissolves during the observation period, by the spouses' education, number of children, and cohort. Qualifying partnerships include legal marriages and cohabiting unions and are defined in Subsection 2.1. The first education listed is the man's and the second is the woman's; $n = 3$ denotes three or more children.

Table A.4: Estimated female Pareto weights by couple type and fertility type

Panel A: Cohort 1967–69

Low fertility type ($r = L$)

	Female education			
Male	Prim.	Sec.	Bach.	MA/PhD
Primary	0.618	0.683	0.681	0.703
Secondary	0.552	0.623	0.627	0.654
Bachelor	0.528	0.603	0.601	0.633
MA/PhD	0.550	0.608	0.604	0.633

High fertility type ($r = H$)

	Female education			
Male	Prim.	Sec.	Bach.	MA/PhD
Primary	0.573	0.674	0.682	0.707
Secondary	0.531	0.616	0.630	0.675
Bachelor	0.510	0.610	0.613	0.657
MA/PhD	0.541	0.608	0.612	0.647

Panel B: Cohort 1978–80

Low fertility type ($r = L$)

	Female education			
Male	Prim.	Sec.	Bach.	MA/PhD
Primary	0.612	0.702	0.703	0.746
Secondary	0.520	0.633	0.643	0.674
Bachelor	0.499	0.608	0.616	0.652
MA/PhD	0.477	0.602	0.610	0.643

High fertility type ($r = H$)

	Female education			
Male	Prim.	Sec.	Bach.	MA/PhD
Primary	0.607	0.729	0.747	0.783
Secondary	0.530	0.662	0.686	0.731
Bachelor	0.507	0.645	0.662	0.708
MA/PhD	0.480	0.634	0.648	0.687

Notes: The table reports estimated female Pareto weights λ^{xy} at the model equilibrium, by couple type (rows: husband education; columns: wife education) and unobserved female fertility type, averaged across locations. In the later cohort, the weight rises with the wife's education and falls with the husband's within every row and column. The weight rises between cohorts in 23 of the 24 cells for wives with at least secondary education, though it falls in some cells for primary-educated wives. The matching-weighted average rises from 0.619 to 0.656; all estimated weights remain interior. The corresponding couple-type weights, averaged over fertility types with matched-mass weights, are plotted in Figure A.9.

Table A.5: Average number of children by couple type: data vs. model

	Cohort 1967–69		Cohort 1978–80	
Couple type (Male, Female)	Data	Model	Data	Model
Primary, Primary	1.882	1.823	1.767	1.504
Primary, Secondary	1.913	1.900	1.849	1.707
Primary, Bachelor	1.909	1.882	1.872	1.743
Primary, Master/PhD	1.649	1.787	1.700	1.730
Secondary, Primary	1.968	1.843	1.875	1.599
Secondary, Secondary	1.934	1.900	1.892	1.782
Secondary, Bachelor	2.000	1.893	1.997	1.820
Secondary, Master/PhD	1.780	1.827	1.836	1.806
Bachelor, Primary	1.878	1.875	1.750	1.659
Bachelor, Secondary	1.882	1.959	1.869	1.839
Bachelor, Bachelor	2.030	1.942	2.023	1.857
Bachelor, Master/PhD	1.874	1.871	1.918	1.846
Master/PhD, Primary	1.837	1.983	1.634	1.691
Master/PhD, Secondary	1.863	1.997	1.809	1.905
Master/PhD, Bachelor	2.028	1.979	1.992	1.902
Master/PhD, Master/PhD	1.956	1.889	1.976	1.875

Notes: The table reports the average number of children by couple education type (the number of children is capped at three), in the data and model, aggregated over fertility types with matched-mass weights and averaged across locations. Couple types are grouped by the male's education.

Table A.6: Matching patterns: data vs. model

	Cohort 1967–69		Cohort 1978–80	
Couple/single type	Data	Model	Data	Model
Panel A: Matched couples (Male, Female)				
Primary, Primary	0.038	0.036	0.023	0.023
Primary, Secondary	0.074	0.069	0.047	0.052
Primary, Bachelor	0.019	0.019	0.019	0.021
Primary, Master/PhD	0.003	0.003	0.004	0.005
Secondary, Primary	0.067	0.057	0.039	0.033
Secondary, Secondary	0.332	0.357	0.263	0.293
Secondary, Bachelor	0.117	0.126	0.151	0.146
Secondary, Master/PhD	0.024	0.024	0.043	0.037
Bachelor, Primary	0.006	0.006	0.004	0.004
Bachelor, Secondary	0.047	0.057	0.040	0.049
Bachelor, Bachelor	0.055	0.057	0.070	0.076
Bachelor, Master/PhD	0.018	0.013	0.035	0.024
Master/PhD, Primary	0.002	0.003	0.002	0.002
Master/PhD, Secondary	0.022	0.021	0.021	0.023
Master/PhD, Bachelor	0.032	0.027	0.046	0.042
Master/PhD, Master/PhD	0.046	0.028	0.085	0.066
Panel B: Single females, by education				
Primary	0.022	0.015	0.018	0.018
Secondary	0.042	0.038	0.041	0.036
Bachelor	0.023	0.028	0.030	0.032
Master/PhD	0.012	0.017	0.021	0.018
Panel C: Single males, by education				
Primary	0.042	0.041	0.041	0.042
Secondary	0.059	0.062	0.061	0.058
Bachelor	0.012	0.011	0.016	0.014
Master/PhD	0.010	0.009	0.014	0.013

Notes: Matched-couple cells in Panel A are shares of the total female population by couple education type, summed over fertility types and averaged across locations. Single shares in Panels B and C are the corresponding female- and male-side population shares; the male shares are expressed in female-population units via the gender ratio, as in the estimation moments.

Table A.7: Female employment rate by education and number of children: data vs. model

	Primary		Secondary		Bachelor		Master/PhD	
	Data	Model	Data	Model	Data	Model	Data	Model
Panel A: Cohort 1967–69, married women								
$n = 0$	0.748	0.877	0.916	0.904	0.938	0.911	0.932	0.949
$n = 1$	0.705	0.801	0.883	0.881	0.936	0.903	0.923	0.965
$n = 2$	0.760	0.784	0.910	0.881	0.953	0.903	0.942	0.957
$n = 3$	0.623	0.668	0.841	0.816	0.930	0.837	0.931	0.912
Panel B: Cohort 1967–69, single women								
$n = 0$	0.477	0.882	0.880	0.909	0.937	0.922	0.935	0.952
$n = 1$	0.493	0.818	0.782	0.843	0.884	0.859	0.812	0.906
$n = 2$	0.491	0.798	0.781	0.817	0.918	0.836	0.797	0.886
$n = 3$	0.410	0.769	0.699	0.786	0.772	0.806	0.772	0.864
Panel C: Cohort 1978–80, married women								
$n = 0$	0.543	0.808	0.843	0.836	0.890	0.872	0.891	0.901
$n = 1$	0.471	0.708	0.815	0.799	0.895	0.868	0.878	0.943
$n = 2$	0.563	0.622	0.876	0.802	0.934	0.870	0.921	0.946
$n = 3$	0.454	0.403	0.822	0.732	0.922	0.838	0.928	0.933
Panel D: Cohort 1978–80, single women								
$n = 0$	0.352	0.881	0.789	0.915	0.884	0.937	0.844	0.962
$n = 1$	0.289	0.612	0.670	0.795	0.827	0.856	0.822	0.906
$n = 2$	0.265	0.353	0.617	0.723	0.821	0.819	0.753	0.880
$n = 3$	0.242	0.150	0.586	0.626	0.728	0.779	0.762	0.851

Notes: The table reports the female employment rate, defined as the probability of working at any intensity, by education, marital status, number of children, and cohort. Here $n = 3$ denotes three or more children. Data moments are computed from the Danish administrative registers; model moments are simulated at the estimated parameters. Values are averaged across locations.

Table A.8: Female full-time share among the employed by education and number of children: data vs. model

	Primary		Secondary		Bachelor		Master/PhD	
	Data	Model	Data	Model	Data	Model	Data	Model
Panel A: Cohort 1967–69, married women								
$n = 0$	0.724	0.805	0.848	0.815	0.814	0.811	0.829	0.838
$n = 1$	0.756	0.723	0.803	0.765	0.795	0.780	0.789	0.835
$n = 2$	0.758	0.750	0.787	0.781	0.785	0.797	0.811	0.835
$n = 3$	0.745	0.712	0.762	0.746	0.747	0.766	0.785	0.805
Panel B: Cohort 1967–69, single women								
$n = 0$	0.626	0.805	0.842	0.808	0.817	0.809	0.822	0.812
$n = 1$	0.754	0.733	0.799	0.736	0.723	0.738	0.810	0.746
$n = 2$	0.765	0.703	0.778	0.706	0.687	0.709	0.745	0.719
$n = 3$	0.791	0.677	0.780	0.680	0.772	0.684	0.828	0.695
Panel C: Cohort 1978–80, married women								
$n = 0$	0.614	0.778	0.817	0.787	0.853	0.805	0.927	0.806
$n = 1$	0.710	0.739	0.824	0.791	0.858	0.824	0.911	0.866
$n = 2$	0.733	0.765	0.832	0.816	0.861	0.844	0.916	0.877
$n = 3$	0.703	0.724	0.790	0.795	0.835	0.836	0.901	0.873
Panel D: Cohort 1978–80, single women								
$n = 0$	0.513	0.845	0.811	0.848	0.858	0.850	0.904	0.852
$n = 1$	0.678	0.718	0.802	0.758	0.822	0.770	0.886	0.778
$n = 2$	0.727	0.623	0.821	0.712	0.793	0.736	0.908	0.746
$n = 3$	0.733	0.524	0.700	0.669	0.852	0.706	0.884	0.716

Notes: The table reports the female full-time share, defined as the probability of working full-time conditional on working, by education, marital status, number of children, and cohort; it complements the employment rates in Table A.7. Here $n = 3$ denotes three or more children. Data moments are computed from the Danish administrative registers; model moments are simulated at the estimated parameters. Values are averaged across locations.

Table A.9: Male employment rate by education and number of children: data vs. model

	Primary		Secondary		Bachelor		Master/PhD	
	Data	Model	Data	Model	Data	Model	Data	Model
Panel A: Cohort 1967–69, married men								
$n = 0$	0.852	0.868	0.908	0.874	0.943	0.932	0.946	0.942
$n = 1$	0.850	0.810	0.918	0.866	0.960	0.947	0.959	0.977
$n = 2$	0.890	0.809	0.931	0.872	0.972	0.955	0.969	0.983
$n = 3$	0.861	0.792	0.880	0.863	0.964	0.954	0.951	0.984
Panel B: Cohort 1967–69, single men								
	0.687	0.917	0.850	0.944	0.939	0.964	0.955	0.975
Panel C: Cohort 1978–80, married men								
$n = 0$	0.696	0.878	0.881	0.865	0.904	0.917	0.908	0.936
$n = 1$	0.708	0.856	0.904	0.910	0.953	0.954	0.940	0.983
$n = 2$	0.791	0.878	0.936	0.922	0.966	0.961	0.967	0.988
$n = 3$	0.731	0.858	0.920	0.904	0.967	0.951	0.974	0.984
Panel D: Cohort 1978–80, single men								
	0.500	0.952	0.801	0.971	0.927	0.980	0.948	0.989

Notes: The table reports the male employment rate, defined as the probability of working at any intensity, by education, marital status, and cohort. Data moments are computed from the Danish administrative registers; model moments are simulated at the estimated parameters. Married rows are split by number of children; $n = 3$ denotes three or more children. Values are averaged across locations.

Table A.10: Male full-time share among the employed by education and number of children: data vs. model

	Primary		Secondary		Bachelor		Master/PhD	
	Data	Model	Data	Model	Data	Model	Data	Model
Panel A: Cohort 1967–69, married men								
$n = 0$	0.790	0.833	0.844	0.805	0.858	0.851	0.862	0.833
$n = 1$	0.786	0.760	0.832	0.783	0.851	0.845	0.856	0.861
$n = 2$	0.784	0.785	0.839	0.807	0.863	0.868	0.863	0.883
$n = 3$	0.769	0.779	0.821	0.806	0.853	0.873	0.877	0.894
Panel B: Cohort 1967–69, single men								
	0.782	0.848	0.825	0.840	0.832	0.829	0.872	0.812
Panel C: Cohort 1978–80, married men								
$n = 0$	0.732	0.844	0.865	0.795	0.869	0.846	0.946	0.834
$n = 1$	0.786	0.824	0.879	0.863	0.922	0.897	0.952	0.910
$n = 2$	0.798	0.870	0.896	0.894	0.936	0.919	0.970	0.931
$n = 3$	0.785	0.872	0.891	0.895	0.937	0.922	0.970	0.933
Panel D: Cohort 1978–80, single men								
	0.783	0.900	0.862	0.892	0.914	0.887	0.969	0.868

Notes: The table reports the male full-time share, defined as the probability of working full-time conditional on working, by education, marital status, and cohort; it complements the employment rates in Table A.9. Data moments are computed from the Danish administrative registers; model moments are simulated at the estimated parameters. Married rows are split by number of children; $n = 3$ denotes three or more children. Values are averaged across locations.

Table A.11: Fertility responses to wages and taxes: external comparison

		Jakobsen et al. (forthcoming)		This paper	
		Data (four-year)	Model (four-year)	Tax reform	Wages +5%
A. Fertility elasticity with respect to own net-of-tax wage					
Women	less skilled	−0.028	−0.021	−0.049	−0.065
	high skilled	−0.011	−0.012	−0.008	−0.034
Men	less skilled	0.023	0.005	0.012	0.027
	high skilled	−0.015	0.002	0.008	0.026
B. Percent change in completed fertility, less-skilled couples, 5 percent wage increase					
Women's wages		—	−1.10	—	−0.32
Men's wages		—	0.30	—	0.14

Notes: The external figures are from the December 2025 version of Jakobsen et al. (forthcoming). Panel A reproduces their Table 4. Its data estimates use tax-induced variation in marginal net wages and fertility over four years; its model column simulates a permanent wage increase from age 35 and measures fertility four years later. Our tax-reform and wage columns measure fertility through age 38 for existing couples in the later-cohort equilibrium, holding their Pareto weights and matching distribution fixed. The wage elasticity divides the percentage change in fertility by the 5 percent wage increase. Less skilled means less than a bachelor's degree, classified by the wife's education in our model. Panel B compares changes in completed fertility after the respective wage increases. Their model applies the wage change from age 25 and allows birth timing to adjust; our model does not choose birth timing. The policy experiments are therefore not identical across studies.

Table A.12: Education distribution by fertility type: sorting is modest and stable

Pr(education | type)	Primary	Secondary	Bachelor	Master's/PhD	Dissimilarity
Low-fertility type, earlier	0.13	0.52	0.26	0.10	0.04
High-fertility type, earlier	0.11	0.56	0.26	0.08	
Low-fertility type, later	0.11	0.44	0.30	0.15	0.04
High-fertility type, later	0.06	0.46	0.33	0.15	

Notes: Model-implied education distribution of each fertility type, by cohort (a model object, since type is latent). The dissimilarity index is $\frac{1}{2}\sum_e |\Pr(e \mid \text{high}) - \Pr(e \mid \text{low})|$, the share of one type that would have to change education to match the other's distribution. Sorting by fertility preference into education is modest and essentially unchanged as education expands, with both types moving up the ladder in parallel.

Table A.13: Policy fertility effects: fixed-couple benchmark, bargaining calculation, and full equilibrium

	Change in fertility, all women			Fixed-couple share
Policy	Fixed couples	+ Bargaining	Full equilibrium	of the full effect
Family transfers				
Free childcare (unconditional)	+0.031	+0.029	+0.054	0.57
Free childcare (mother works full-time)	+0.023	+0.021	+0.040	0.57
Taxing the couple				
Joint taxation (income aggregation)	−0.011	−0.012	−0.037	0.31
Wage policies				
Remove motherhood wage penalty	+0.016	+0.016	+0.027	0.60
Close gender wage gap	−0.013	−0.020	−0.037	0.36
Education subsidies				
Men's education subsidy	+0.000	−0.002	+0.015	0.00
Women's education subsidy	+0.000	+0.000	+0.014	0.00

Notes: Changes in completed fertility per woman from the later-cohort baseline. "Fixed couples" applies the policy to household budgets and lets baseline couples re-optimize labor supply and fertility at baseline Pareto weights. "+ Bargaining" additionally imposes the Pareto weights from the policy equilibrium while retaining baseline couple masses; this intermediate calculation is sequence-dependent and is not a separate equilibrium. "Full equilibrium" jointly re-solves education, partnership formation, partner composition, and Pareto weights. The final column is the ratio of the fixed-couple effect to the full-equilibrium effect. Education subsidies change no household budget, so their fixed-couple effect is zero by construction.

Table A.14: Upstream responses to each policy: education, matching, and partnership

	BA or above (pp)		Matching		Partnership
Policy	Women	Men	Educ. corr.	Same education (pp)	rate (pp)
Family transfers					
Free childcare (unconditional)	+0.8	+1.2	−0.004	+0.2	+1.5
Free childcare (mother works full-time)	+1.5	+1.4	−0.001	+0.1	+1.1
Taxing the couple					
Joint taxation (income aggregation)	−4.5	−5.7	−0.032	+0.2	−1.6
Wage policies					
Remove motherhood wage penalty	−0.3	−0.6	−0.015	+0.2	+0.6
Close gender wage gap	+1.3	+0.4	−0.003	−0.2	−1.1
Education subsidies					
Men's education subsidy	+11.4	+28.1	+0.006	−0.1	+0.5
Women's education subsidy	+29.7	+15.4	−0.040	−6.8	+0.0

Notes: Changes from the later-cohort baseline, re-solving the full equilibrium. The first two columns give the change in the share of women and of men who choose a bachelor's degree or above; "Educ. corr." is the husband–wife education correlation and "Same education" the share of couples in which both spouses hold the same level; the last column is the economy-wide partnership rate. Baseline values: 0.467 and 0.305 for the education shares, 0.445 for the correlation, 0.511 for the same-education share, and 0.896 for the partnership rate.

Table A.15: Joint taxation: responses by couple type

Husband	Wife	Share of couples	Δ Fertility	Δ Both full-time
Primary	Primary	0.026	+0.002	-0.074
	Secondary	0.059	+0.000	-0.131
	Bachelor	0.023	+0.000	-0.158
	MA/PhD	0.006	-0.013	-0.162
Secondary	Primary	0.037	-0.009	-0.085
	Secondary	0.327	-0.010	-0.141
	Bachelor	0.163	-0.012	-0.165
	MA/PhD	0.042	-0.018	-0.169
Bachelor	Primary	0.005	-0.017	-0.101
	Secondary	0.054	-0.015	-0.149
	Bachelor	0.085	-0.017	-0.162
	MA/PhD	0.027	-0.024	-0.158
MA/PhD	Primary	0.002	-0.017	-0.100
	Secondary	0.025	-0.017	-0.135
	Bachelor	0.047	-0.020	-0.142
	MA/PhD	0.074	-0.028	-0.135

Notes: Each row reports the change from the later-cohort baseline under joint taxation for one couple type, defined by the husband's and the wife's education. "Share of couples" is the cell's share of all couples at the baseline. "Δ Fertility" is the change in expected children per couple; "Δ Both full-time" is the change in the probability that both spouses work full-time, measured as a fraction rather than in percentage points. Cells are averaged across locations and fertility types, holding the baseline matched masses fixed. The entries therefore describe responses within couple types rather than changes in their partnership rates.

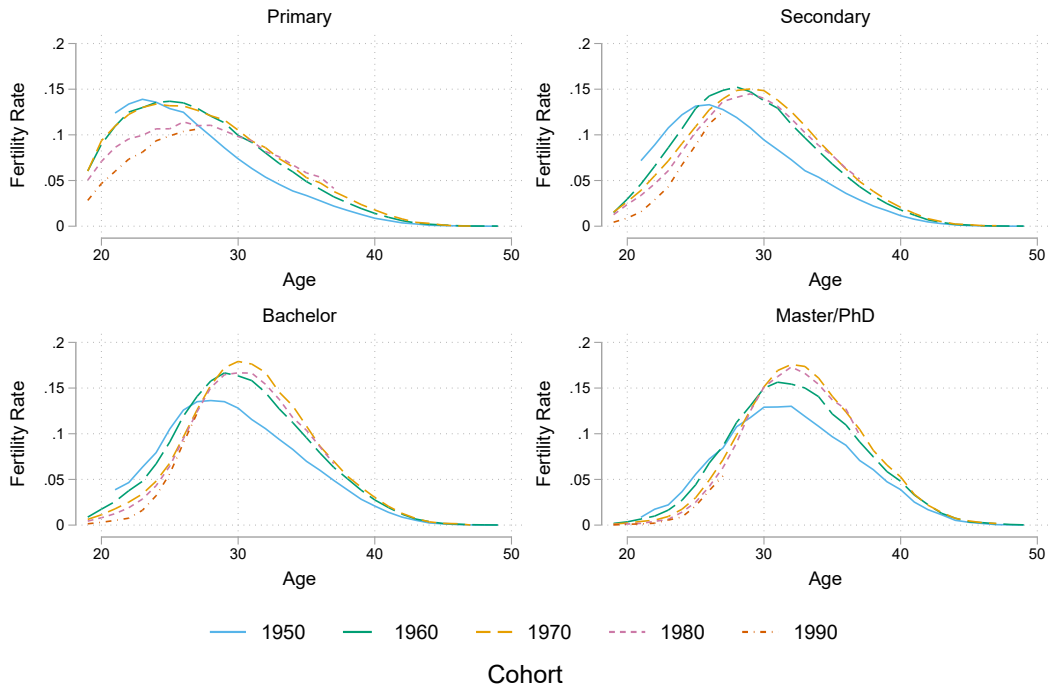
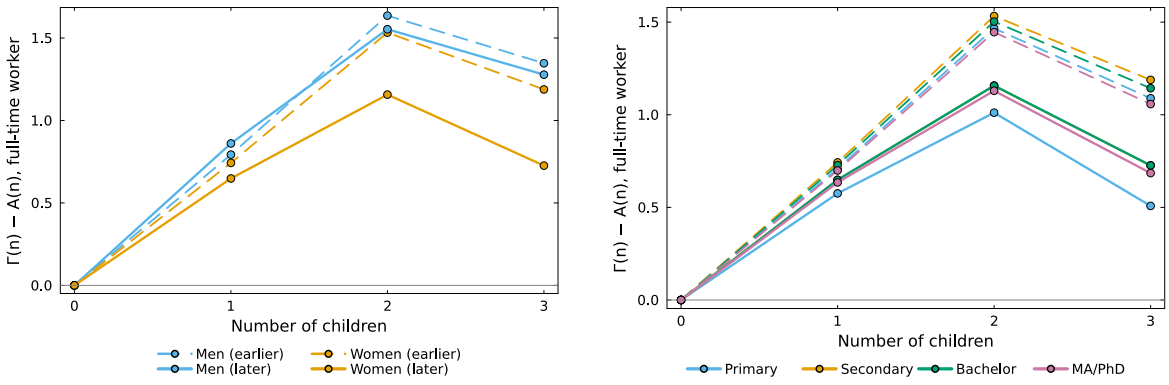


Figure A.1: Age-specific fertility rate across age, cohort, and education

Notes: Data are from Danish administrative registers provided by Statistics Denmark. The figure reports age-specific fertility rates, defined as births per woman at each age, by education for the birth-cohort groups labeled 1950, 1960, 1970, 1980, and 1990. For cohorts observed beyond age 38, fertility rates are low after that age. This pattern supports measuring cohort fertility over the common age window from 19 through 38 in Section 2.

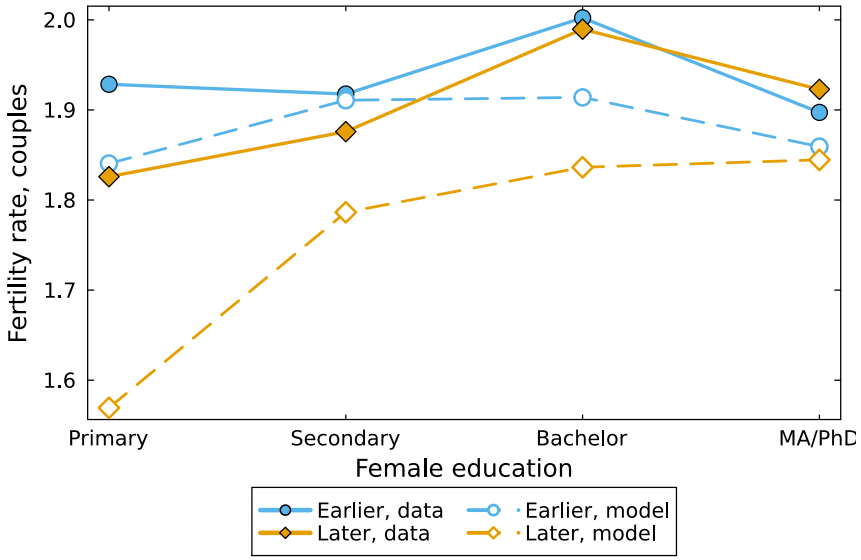
Figure A.2: Utility from children net of full-time work–family costs



(a) Men and secondary-educated women (b) Women by education

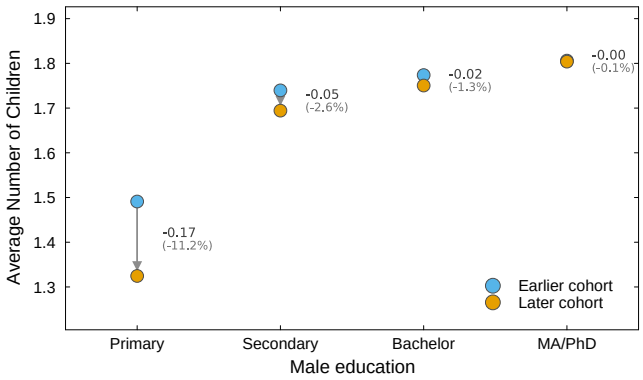
Notes: The figure plots utility from children net of the work–family cost of full-time work, $\Gamma_{g,r}(n) - \alpha_g(n;e)$, using the work–family cost for partnered women. Female preferences are averaged across fertility types using population shares. Panel (a) compares men with secondary-educated women; panel (b) shows women in all four education groups. Dashed lines denote the earlier cohort and solid lines the later cohort.

Figure A.3: Couples fertility rate by female education: data vs. model



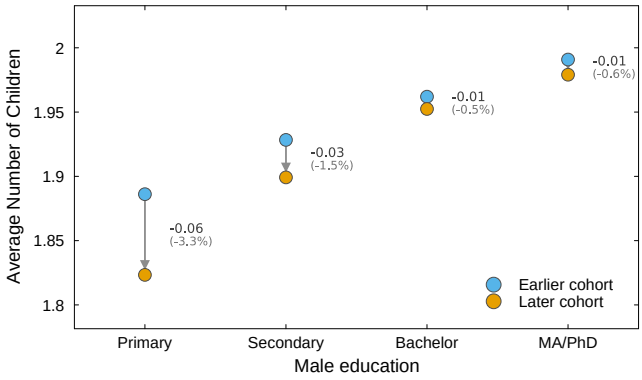
Notes: Data are from Danish administrative registers provided by Statistics Denmark; model rates use the estimated parameters. The earlier cohort comprises women born in 1967–69 and the later cohort women born in 1978–80. Fertility is the average number of children per woman accumulated between ages 19 and 38, conditional on having a qualifying legal marriage or cohabiting partnership as defined in Subsection 2.1, and is reported by highest education completed by age 38. Filled markers with solid lines denote data and open markers with dashed lines model predictions; light blue denotes the earlier cohort and orange the later cohort.

Figure A.4: Average number of children per man, by male education and cohort



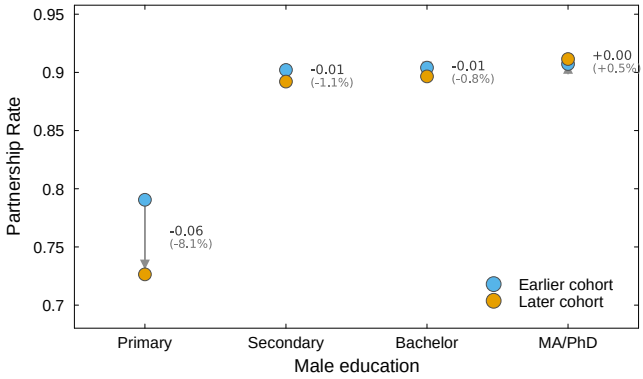
Notes: Data are from Danish administrative registers provided by Statistics Denmark. Men born in 1967–69 form the earlier cohort, and men born in 1978–80 form the later cohort. Fertility is the average number of children per man accumulated between ages 19 and 38. The figure reports fertility by male education and cohort. Each pair of points reports the earlier- and later-cohort values for an education group, with the across-cohort change reported alongside.

Figure A.5: Average number of children per man *conditional on being in a couple*, by male education and cohort



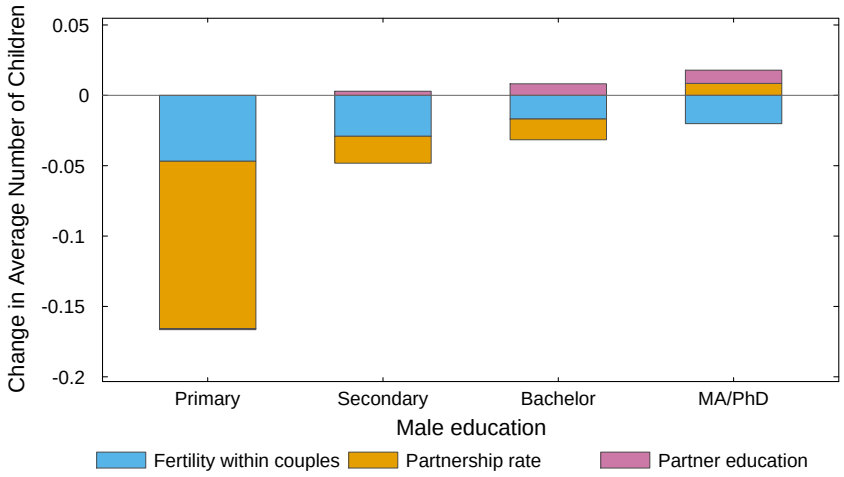
Notes: Data are from Danish administrative registers provided by Statistics Denmark. Men born in 1967–69 form the earlier cohort, and men born in 1978–80 form the later cohort. Fertility is the average number of children per man accumulated between ages 19 and 38. The figure reports fertility conditional on being in a couple, by male education and cohort. Each pair of points reports the earlier- and later-cohort values for an education group, with the across-cohort change reported alongside.

Figure A.6: Partnership rate, by male education and cohort



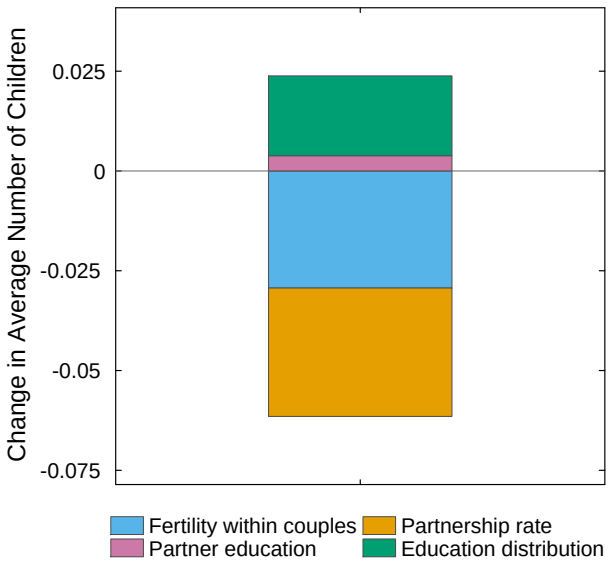
Notes: Data are from Danish administrative registers provided by Statistics Denmark. Men born in 1967–69 form the earlier cohort, and men born in 1978–80 form the later cohort. The partnership rate is the share of men classified as partnered under the definition in Subsection 2.1. The figure reports this rate by male education and cohort. Each pair of points reports the earlier- and later-cohort values for an education group, with the across-cohort change reported alongside.

Figure A.7: Shapley decomposition of the fertility change, by male education



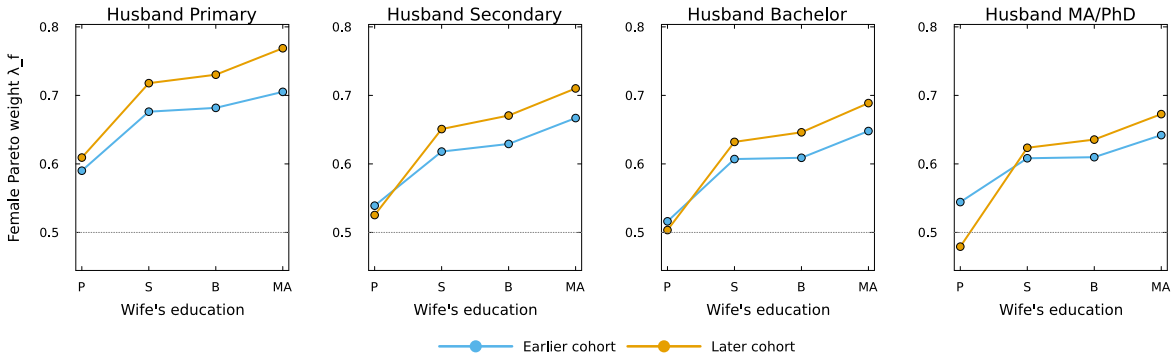
Notes: Data are from Danish administrative registers provided by Statistics Denmark. Men born in 1967–69 form the earlier cohort, and men born in 1978–80 form the later cohort. Fertility is the average number of children per man accumulated between ages 19 and 38. Within each male education group, the figure decomposes the across-cohort change into fertility within couples, the partnership rate, and partner education. Contributions are measured in children per man. Each Shapley value averages a component's marginal contribution over all six possible replacement orders, and the contributions sum to the total change within each education group. The construction reverses the roles of women and men in Figure 4.

Figure A.8: Shapley decomposition of the economy-wide fertility change, men



Notes: Data are from Danish administrative registers provided by Statistics Denmark. Men born in 1967–69 form the earlier cohort, and men born in 1978–80 form the later cohort. Fertility is the average number of children per man accumulated between ages 19 and 38. The figure decomposes the economy-wide change into fertility within couples, the partnership rate, partner education, and the male education distribution. Contributions are measured in children per man. Each Shapley value averages a component's marginal contribution over all 24 possible replacement orders, and the contributions sum to the economy-wide change. The construction reverses the roles of women and men in Figure 5.

Figure A.9: Wife's Pareto weight by couple type



Notes: Equilibrium Pareto weight on the wife, λ^{xy}, for each couple type, averaged over fertility types and locations with the matched masses as weights. Each panel fixes the husband's education and varies the wife's; blue denotes the earlier cohort and orange the later. A higher value places more weight on the wife's preferences in the household problem. Weights are read from the estimated equilibrium of each cohort; Table A.4 reports them by fertility type.

B Decomposition Methods

We use Shapley values in three across-cohort decompositions: one based on the data, one based on changes in parameters and inputs, and one that separates household primitives, Pareto weights, and population composition. Starting from an initial outcome, a Shapley value averages a component's incremental contribution over all orders in which the components can be replaced by their final values. The same rule decomposes the fertility effect of free childcare. The fixed-couple policy comparison instead changes margins in a specified order.

B.1 Descriptive decomposition of the data (Section 2)

Section 2 divides the across-cohort change in fertility among fertility within couples, partnership rates, partner education conditional on partnership, and the female education distribution. Within a female education group, only the first three components vary.

B.1.1 Setup and notation

Let $x, y \in \{P, S, B, M\}$ index female and male education, respectively, where P, S, B, and M denote primary, secondary, bachelor's, and master's/PhD education. For a given cohort, let

- M_{xy} denote the mass of couples in which the woman has education x and the man has education y;
- S_x^F denote the mass of women with education x who are classified as single;
- $M_x \equiv \sum_y M_{xy}$ denote the mass of partnered women with education x;
- $W_x \equiv M_x + S_x^F$ denote the total mass of women with education x;
- n_{xy} denote mean fertility among couples in which the woman has education x and the man has education y.

For each cohort, define the female education share, partnership rate, and partner-education distribution as follows:

$$W_x^{\text{share}} \equiv \frac{W_x}{\sum_{x'} W_{x'}} \qquad \text{(female education distribution)}, \tag{B.1}$$

$$\chi_x \equiv \Pr(\text{partnered} \mid x) = \frac{M_x}{M_x + S_x^F} \qquad \text{(partnership rate at female education } x\text{)}, \tag{B.2}$$

$$\varpi_{y|x} \equiv \Pr(y \mid x, \text{ partnered}) = \frac{M_{xy}}{\sum_{y'} M_{xy'}} \qquad \text{(partner-education distribution at } x\text{)}. \tag{B.3}$$

Note that $\sum_y \varpi_{y|x} = 1$ for every x and $\sum_x W_x^{\text{share}} = 1$.

In the decomposition, let r_x^w denote average fertility among women with education x, and

let R^w denote average fertility among all women:

$$r_x^w = \chi_x \sum_y \varpi_{y|x} n_{xy}, \tag{B.4}$$

$$R^w = \sum_x W_x^{\text{share}} \cdot r_x^w = \sum_x W_x^{\text{share}} \chi_x \sum_y \varpi_{y|x} n_{xy}. \tag{B.5}$$

The sum $\sum_y \varpi_{y|x} n_{xy}$ is mean fertility among partnered women with education x. Equation (B.5) averages the education-specific rates using the female education distribution.

B.1.2 Shapley decomposition by female education

We decompose the cohort change in r_x^w into the contributions of fertility within couples, the partnership rate, and the conditional partner-education distribution. Let $K_x = \{n, \chi, \varpi\}$ denote these three components. For each subset $A \subseteq K_x$, let $v_x(A)$ be the fertility rate obtained from (B.4) when the components in A take their later-cohort values and the remaining components take their earlier-cohort values. Thus,

$$v_x(\varnothing) = \chi_x^{\text{earlier}} \sum_y \varpi_{y|x}^{\text{earlier}} n_{xy}^{\text{earlier}} = r_x^{w,\text{earlier}}, \tag{B.6}$$

$$v_x(K_x) = \chi_x^{\text{later}} \sum_y \varpi_{y|x}^{\text{later}} n_{xy}^{\text{later}} = r_x^{w,\text{later}}. \tag{B.7}$$

For component $k \in K_x$, its Shapley contribution is

$$\phi_{xk} = \sum_{A \subseteq K_x \setminus \{k\}} \frac{|A|!\big(|K_x| - |A| - 1\big)!}{|K_x|!} \left[v_x(A \cup \{k\}) - v_x(A) \right]. \tag{B.8}$$

Equation (B.8) averages the marginal contribution of component k over all $3! = 6$ possible replacement orders. The factorial term gives each replacement order equal weight. The three contributions exhaust the cohort change:

$$\Delta r_x^w = r_x^{w,\text{later}} - r_x^{w,\text{earlier}} = \phi_{xn} + \phi_{x\chi} + \phi_{x\varpi}. \tag{B.9}$$

B.1.3 Shapley decomposition of the economy-wide rate

The economy-wide decomposition adds the female education distribution to the three components of the education-specific decomposition. Let $K = \{W, n, \chi, \varpi\}$ denote the four components in (B.5). For each subset $A \subseteq K$, let $V(A)$ be the economy-wide fertility rate obtained when the components in A take their later-cohort values and the remaining components take their earlier-cohort values. In particular, $V(\varnothing) = R^{w,\text{earlier}}$ and $V(K) = R^{w,\text{later}}$. The Shapley contribution of component $k \in K$ is

$$\Phi_k = \sum_{A \subseteq K \setminus \{k\}} \frac{|A|!\big(|K| - |A| - 1\big)!}{|K|!} \left[V(A \cup \{k\}) - V(A) \right]. \tag{B.10}$$

Equation (B.10) averages each component's marginal contribution over all $4! = 24$ possible replacement orders. The four contributions sum to the economy-wide cohort change:

$$\Delta R^w = R^{w,\text{later}} - R^{w,\text{earlier}} = \Phi_W + \Phi_n + \Phi_\chi + \Phi_\varpi. \tag{B.11}$$

B.1.4 Interpretation and limitations

The education shares, partnership rates, partner-education distributions, and couple fertility used here all reflect underlying preferences and market conditions. A change in partner education may reflect changes in matching surplus or in the education distributions of men and women. A change in n_{xy} may reflect preferences, wages, work–family costs, bargaining, or selection into education and marriage. The descriptive decomposition therefore assigns parts of the fertility change to observed margins without identifying the forces that moved them.

B.2 Structural Shapley decomposition by parameter block (Section 6.3)

We attribute the fertility difference between the two estimated cohort equilibria to nine disjoint blocks of parameters and inputs. They are men's and women's utility from children, men's and women's work–family costs, men's and women's wages, match value, the high-fertility-type share, and all remaining changes. Let $\mathcal{B}$ denote these blocks. For each subset $A \subseteq \mathcal{B}$, we assign later-cohort values to the blocks in A, keep the other blocks at their earlier-cohort values, and solve for equilibrium. Let $G(A)$ denote the resulting economy-wide fertility rate. Thus, $G(\varnothing)$ is the earlier-cohort model outcome and $G(\mathcal{B})$ is the later-cohort outcome. The Shapley contribution of block $b \in \mathcal{B}$ is

$$\Psi_b = \sum_{A \subseteq \mathcal{B}\setminus\{b\}} \frac{|A|!\big(|\mathcal{B}| - |A| - 1\big)!}{|\mathcal{B}|!} \left[G(A \cup \{b\}) - G(A)\right]. \tag{B.12}$$

Computing (B.12) requires $2^9 = 512$ equilibrium solutions. The block contributions allocate all interactions and satisfy

$$G(\mathcal{B}) - G(\varnothing) = \sum_{b \in \mathcal{B}} \Psi_b. \tag{B.13}$$

We further divide each Ψ_b into three outcome components: the female education distribution, partnership and matching, and fertility within couple types. Define

$$q_{y|x} \equiv \chi_x \varpi_{y|x} = \frac{M_{xy}}{W_x}, \tag{B.14}$$

the probability that a woman with education x is partnered with a man of education y. Because $\sum_y q_{y|x} = \chi_x$, a change in $q_{y|x}$ captures both the partnership rate and partner education conditional on partnership. Equation (B.5) can then be written as

$$R^w = \sum_x W_x^{\text{share}} \sum_y q_{y|x} n_{xy}. \tag{B.15}$$

For each equilibrium solution, collect the three components in $z_A = (W_A, q_A, n_A)$. We use the

Shapley formula to divide the fertility change from A to $A \cup \{b\}$ among these components. Let $d_{bk}(A)$ denote the resulting contribution of outcome component $k \in K^S = \{W, q, n\}$. For $T \subseteq K^S$, define $h_{A,b}(T)$ by using the components of $z_{A\cup\{b\}}$ in T and the components of z_A outside T in (B.15). Then

$$d_{bk}(A) = \sum_{T \subseteq K^S \setminus \{k\}} \frac{|T|!(|K^S| - |T| - 1)!}{|K^S|!} \big[h_{A,b}(T \cup \{k\}) - h_{A,b}(T)\big]. \tag{B.16}$$

Averaging these outcome contributions with the parameter-block Shapley weights gives

$$\Psi_{bk} = \sum_{A \subseteq \mathcal{B} \setminus \{b\}} \frac{|A|!\big(|\mathcal{B}| - |A| - 1\big)!}{|\mathcal{B}|!} d_{bk}(A). \tag{B.17}$$

For every block, $\sum_{k \in K^S} \Psi_{bk} = \Psi_b$. For each one-block change, the three contributions sum to $G(A \cup \{b\}) - G(A)$. Each row of Table B.1 sums to its block's Shapley value; summing either the rows or the columns gives the fertility difference between the two estimated cohort equilibria. The fertility-within-couple-types component changes only n_{xy}, holding the education distribution and the entire distribution of matches fixed. The partnership-and-matching component changes $q_{y|x}$ and therefore includes both whether individuals form partnerships and who partners with whom. Both decompositions use the fertility identity in (B.5), but this calculation groups the partnership rate and partner education and re-solves the model for each set of parameter blocks.

B.3 Structural Shapley decomposition by margin (Section 6)

Changes in Pareto weights enter the parameter-block contributions but have no separate entry in that decomposition. We therefore use a second Shapley calculation to separate household primitives, Pareto weights, and the masses of coupled and single individuals. Let $c \in \{0, 1\}$ index the earlier and later cohorts, and let $\boldsymbol{\Theta}_c$ collect the cohort's household primitives and exogenous inputs. Solving the model yields Pareto weights $\boldsymbol{\lambda}_c = \{\lambda^k_{xyr,c}\}$ and masses $\{M^k_{xyr,c}, S^k_{xr,c}\}$ in each local market k. Here x and y index female and male education, and r indexes the woman's fertility type. $M^k_{xyr,c}$ and $S^k_{xr,c}$ are the masses of couples of type (x, y, r) and single women of type (x, r), respectively. Let $n^k_{xyr}(\boldsymbol{\Theta}, \lambda^k_{xyr})$ denote expected fertility for a couple of that type. For any combination of the three objects, define

$$R(\boldsymbol{\Theta}, \boldsymbol{\lambda}, M, S) = \frac{\sum_{k,x,y,r} M^k_{xyr}\, n^k_{xyr}(\boldsymbol{\Theta}, \lambda^k_{xyr})}{\sum_{k,x,y,r} M^k_{xyr} + \sum_{k,x,r} S^k_{xr}}. \tag{B.18}$$

The equilibrium fertility rate for cohort c is $R_c \equiv R(\boldsymbol{\Theta}_c, \boldsymbol{\lambda}_c, M_c, S_c)$.

We replace three arguments of R separately:

- *within-couple*: replace $\boldsymbol{\Theta}_0$ with $\boldsymbol{\Theta}_1$ inside each couple's fertility problem;
- *bargaining*: replace $\boldsymbol{\lambda}_0$ with $\boldsymbol{\lambda}_1$ inside that problem;
- *composition*: replace (M_0, S_0) with (M_1, S_1), thereby changing partnership formation, matching, education, and fertility-type composition.

Table B.1: Decomposition of the across-cohort fertility change by parameter block

		Outcome margin		
Parameter or input block	Shapley value	Within couples	Partnership and matching	Education
Men's utility from children (Γ_m)	−0.001	+0.000	−0.001	−0.000
Women's utility from children ($\Gamma_{f,r}$)	−0.238	−0.119	−0.113	−0.006
Men's work–family cost (α_m)	+0.002	−0.005	+0.007	+0.000
Women's work–family cost (α_f)	+0.031	+0.007	+0.016	+0.008
Men's wages	−0.002	+0.002	−0.004	−0.000
Women's wages	−0.001	+0.001	−0.003	+0.001
Match value (ξ)	+0.156	−0.037	+0.182	+0.011
High-fertility-type share (π_H)	−0.048	−0.041	−0.007	−0.001
Other changes	−0.003	+0.092	−0.081	−0.014
All blocks (later-cohort equilibrium)	−0.105	−0.101	−0.003	−0.001

Notes: Each row reports the Shapley value of a disjoint parameter or input block for the across-cohort change in model-implied fertility per woman. For every subset of the nine blocks, we replace those blocks with their later-cohort values and re-solve the equilibrium; each row averages the block's marginal contribution over all replacement orders. Each preference block changes one gender's full function $\Gamma_g(n)$, including its intercept and curvature; the implementation keeps the two gender-specific changes disjoint although these coefficients are constrained to be equal within each cohort. The wage blocks change cohort-specific wage schedules. "Other" includes labor disutility, shock parameters, divorce risk, and sex ratios. Each block's value is further divided by a three-factor Shapley calculation into fertility within fixed couple types, partnership and matching, and women's education distribution. Rows and columns sum to the model's fertility change apart from rounding. The main-text figure omits the Other bar. Section B.2 gives the construction.

Let $J = \{\Theta, \lambda, C\}$ index within-couple primitives, Pareto weights, and composition, where $C = (M, S)$. For each subset $A \subseteq J$, let $v(A)$ denote economy-wide fertility from (B.18) when the arguments in A take their later-cohort values and the remaining arguments take their earlier-cohort values. We evaluate all $2^3 = 8$ combinations. The Shapley contribution of margin $j \in J$ is

$$\Omega_j = \sum_{A \subseteq J \setminus \{j\}} \frac{|A|!\big(|J| - |A| - 1\big)!}{|J|!} \left[v(A \cup \{j\}) - v(A)\right]. \tag{B.19}$$

Equation (B.19) averages over the six replacement orders, allocating interactions across the three margins. The contributions sum to the fertility difference between the two cohort equilibria:

$$R_1 - R_0 = \Omega_\Theta + \Omega_\lambda + \Omega_C. \tag{B.20}$$

Only $v(\varnothing) = R_0$ and $v(J) = R_1$ are solved cohort equilibria. The other six values combine objects from the two equilibria without re-solving the marriage market.

B.4 Policy benchmarks (Section 7)

The policy comparison uses the same fertility function R, but changes its arguments in a fixed order. Subscript 0 now denotes the later-cohort baseline and subscript 1 the new equilibrium under a policy. Figure 13 compares the first and last calculations below, while Table A.13 also reports the intermediate bargaining calculation. Before subtracting baseline fertility R_0, they

are

$$R^{\text{fixed}} = R(\boldsymbol{\Theta}_1, \boldsymbol{\lambda}_0, M_0, S_0), \tag{B.21}$$

$$R^{\text{barg}} = R(\boldsymbol{\Theta}_1, \boldsymbol{\lambda}_1, M_0, S_0), \tag{B.22}$$

$$R^{\text{full}} = R(\boldsymbol{\Theta}_1, \boldsymbol{\lambda}_1, M_1, S_1) = R_1. \tag{B.23}$$

The fixed-couple calculation lets baseline couples change labor supply and fertility under the policy while holding their Pareto weights and all couple and single masses fixed. The intermediate calculation also applies the policy-equilibrium Pareto weights to baseline couples. The full-equilibrium calculation also uses the policy-equilibrium masses of coupled and single individuals, which reflect changes in education, partnership formation, partner composition, and fertility-type composition. Neither of the first two calculations is a new marriage market equilibrium.

B.5 Shapley decomposition of childcare effects (Section 7.2)

For the childcare counterfactuals, we calculate four Shapley contributions and report them in three groups in Figure 14. Within local market k, let s^k_{xr} be the share of women with education x and fertility type r, and let p^k_{xr} be their partnership rate. Conditional on partnership, $\pi^k_{y|xr}$ is the distribution of partner education; n^k_{xyr} is expected fertility for couple type (x, y, r). The accounting identity is

$$F^k(s, p, \pi, n) = \sum_{x,r} s^k_{xr} p^k_{xr} \sum_y \pi^k_{y|xr} n^k_{xyr}. \tag{B.24}$$

For each factor in $J_C = \{s, p, \pi, n\}$, we replace its baseline value with its value in the new policy equilibrium. Let $H^k(A)$ be the value of (B.24) when factors in $A \subseteq J_C$ take their policy values and the others retain their baseline values. The contribution of factor j in market k is

$$\phi^k_j = \sum_{A \subseteq J_C \setminus \{j\}} \frac{|A|!(|J_C| - |A| - 1)!}{|J_C|!} \left[H^k(A \cup \{j\}) - H^k(A) \right]. \tag{B.25}$$

We calculate the four contributions within each market. Write ϕ_j for the arithmetic mean of ϕ^k_j across markets. The partnership component is ϕ_p, fertility within couple types is ϕ_n, and the figure combines $\phi_s + \phi_\pi$ as education, fertility-type, and partner composition. These contributions sum to the change in the accounting rate F under the policy. The mixed values combine outcomes from two solved equilibria; the model is not re-solved for each replacement order.

C Danish Institutions in the Model

This appendix specifies the tax and transfer functions used in the household budget constraints. We adapt the Danish institutions described by Jakobsen et al. (forthcoming) to the static model. Unless otherwise stated, monetary amounts are annual and measured in DKK 1,000.

C.1 Unemployment benefit: $\mathcal{B}$

We assume that individuals who do not work ($l = 0$) receive unemployment benefits of DKK 118,284 ($\mathcal{B} = 118.284$), which is equal to the annual basic income assistance level in Denmark in 2010. We assume that this value remains unchanged over time and cohort.

C.2 Tax: $\mathcal{T}$

We use a simplified version of Denmark's 2010 tax system, following Jakobsen et al. (forthcoming). Let Y_j and Y_{-j} denote the incomes of individual j and their spouse, respectively. Individual tax liability is calculated as follows:

$$
\begin{aligned}
\tau_{\max} &= \tau_l + \tau_u + \tau_c + \tau_h - \bar{\tau}, \\
\text{personal income} &= (1 - \mathbb{I}(l_j > 0)\tau_{\text{LMC}}) \cdot Y_j, \\
\text{taxable income} &= \text{personal income} - \min\{WD \cdot Y_j, \overline{WD}\}, \\
\underline{y}_l &= \underline{y} + \mathbb{I}(\text{couple}) \cdot \max\{0, \underline{y} - \text{personal income}_{-j}\}, \\
T_c &= \max\{0, \tau_c \cdot (\text{taxable income} - \underline{y})\}, \\
T_h &= \max\{0, \tau_h \cdot (\text{taxable income} - \underline{y})\}, \\
T_l &= \max\{0, \tau_l \cdot (\text{personal income} - \underline{y}_l)\}, \\
T_u &= \max\{0, (\tau_u - \max\{0, \tau_{\max}\}) \cdot (\text{personal income} - \underline{y}_u)\}, \\
\text{tax} &= \mathbb{I}(l_j > 0)\tau_{\text{LMC}} Y_j + T_c + T_h + T_l + T_u
\end{aligned}
$$

Table C.1 reports the parameter values. When the sum of the bottom, top, municipal, and health rates exceeds $\bar{\tau}$, the top-bracket rate is reduced by the excess. At the calibrated values, its effective increment is 0.1419 and the top marginal rate, including the labor market contribution, is 55.38 percent. The municipal rate includes church tax; the model does not distinguish individual municipalities or church membership.

C.3 Child benefits and childcare costs: $\mathcal{C}$

The static model makes childcare fees and child benefits functions of family size, household income, and single-parent status. Let n denote the number of children, Y total household income before tax, and s an indicator for a single-parent household. Net child-related expenditure in the budget constraint is

$$\mathcal{C}(n, Y, s) = C_c(n, Y, s) - C_s(n, s), \tag{C.1}$$

where C_c denotes childcare fees after income-related rebates and C_s denotes cash child benefits. Both are zero when $n = 0$.

For households with children, fees before the income-related rebate are

$$\widetilde{C}_c(n) = a_0 + a_1 + \frac{a_0}{2}(n - 1), \qquad n \geq 1, \tag{C.2}$$

Table C.1: Tax System Parameters in 2010

Symbol	Value in 2010 (DKK)	Description
$\bar{\tau}$	.515	Maximum tax rate
τ_{LMC}	.08	Labor Market Contribution
WD	.0425	EITC rate
$\overline{WD}$	13,600	Maximum EITC
τ_c	.2564	Municipal tax rate (incl. church tax)
$\underline{y}$	42,900	Personal allowance
$\underline{y}_u$	389,900	Top-bracket threshold
τ_h	.08	Health contribution tax
τ_l	.0367	Tax rate in lowest tax bracket
τ_u	.15	Tax rate in upper tax bracket

Notes: The table reports the parameters of the Danish personal income tax system in 2010, the year used to calibrate the tax and transfer functions in the model. Symbols correspond to those used in the budget constraints of Section 3. Monetary amounts are in Danish kroner.

with $a_0 = 6.109$ and $a_1 = 21.127$. The first child is assigned the higher fee of DKK 27,236, and each additional child adds DKK 3,054.50 after the 50-percent sibling rebate. This is a static approximation that assigns the first child to the younger age group and the remaining children to the older group; children's ages are not state variables.[27]

Households pay a fraction $\widetilde{\pi}(n, Y, s)$ of this fee:

$$C_c(n, Y, s) = \widetilde{\pi}(n, Y, s)\widetilde{C}_c(n), \qquad n \geq 1, \tag{C.3}$$

$$\widetilde{\pi}(n, Y, s) = \begin{cases} 0, & Y < Y_{\min}(n, s), \\ \min\{1, \zeta_0 + \zeta_1[Y - Y_{\min}(n, s)]\}, & Y \geq Y_{\min}(n, s), \end{cases} \tag{C.4}$$

$$Y_{\min}(n, s) = 150 + 7(n - 1) + 60s, \tag{C.5}$$

where $\zeta_0 = 0.05$ and $\zeta_1 = 0.0026$. The income threshold for a full exemption increases with the number of children and is higher for single parents. Above this threshold, the fee rises with income until the household pays the full amount.

Cash child benefits are

$$C_s(n, s) = \mathbf{1}\{n > 0\}\left[c_0 + c_1(n - 1) + s(c_2 + c_3 n)\right], \tag{C.6}$$

where $c_0 = 16.988$, $c_1 = 10.580$, $c_2 = 4.956$, and $c_3 = 4.868$. The first two terms assign a higher benefit to the first child and a lower benefit to each additional child, using the same static age approximation as the fee schedule.[28] The final terms provide a fixed benefit and a per-child supplement for single parents. These benefits are tax-exempt and do not depend on household

[27] In the age-dependent formulation of Jakobsen et al. (forthcoming), childcare fees and benefits depend on the number of children n_t and the youngest child's age o_t. Writing $b(o_t)$ for the indicator that the youngest child is in the young-child category, the fee schedule has the form $a_0 + a_1 b(o_t) + (a_0/2)(n_t - 1)$ and ordinary child benefits have the form $c_0 b(o_t) + c_1[n_t - b(o_t)]$. The static implementation sets this indicator to one whenever $n > 0$.

[28] The age-dependent amounts reported by Jakobsen et al. (forthcoming) under the 2010 rules are DKK 16,988 for ages 0–2, DKK 13,448 for ages 3–6, and DKK 10,580 for ages 7–17. The model uses the first and last amounts rather than tracking each child's age.

income.

The free-childcare counterfactuals set C_c to zero for eligible households and leave C_s unchanged. The unconditional policy applies to all households with children; the work-contingent policy applies only when the mother works full-time, including after dissolution. Consequently, free childcare eliminates fees rather than setting net child-related expenditure $\mathcal{C}$ to zero.

C.4 Child support: $\mathcal{D}$

In case of divorce, we assume that the man pays child support to his former wife. We assume that the total annual amount of child support the man has to transfer is thus ζn where $\zeta = 14,040$ DKK in 2010.

D Auxiliary Equilibrium: Numerical Implementation

This appendix describes the algorithm used to compute the competitive auxiliary equilibrium for given model primitives and fertility-type shares. The algorithm is applied separately to each marriage market, with market-specific population masses and type shares. The algorithm extends the iteration process of Gayle and Shephard (2019), which builds on Galichon et al. (2019), by jointly updating the endogenous female and male education-type masses. The partner-choice probabilities incorporate the partner-mass adjustment implied by the match-specific preference shocks of Gutierrez (2020). At this stage of the estimation, fertility types have already been estimated in the first step. The equilibrium solver takes the estimated type shares $\{\pi_{r,k,c}\}_{r\in\mathcal{R}}$ for each market k and cohort c as given; both indices are suppressed in the equations below.

Step 1: initialization. Let k index the outer iteration. Choose strictly positive initial type masses $\{f_x^{(0)}\}_{x\in\mathcal{X}}$ and $\{m_y^{(0)}\}_{y\in\mathcal{Y}}$ satisfying

$$\sum_{e\in\mathcal{E}} f_{(e,r)}^{(0)} = \mu_{\mathcal{X}}\pi_r \quad \text{for every } r\in\mathcal{R}, \qquad \sum_{y\in\mathcal{Y}} m_y^{(0)} = \mu_{\mathcal{Y}},$$

together with strictly positive initial single-choice probabilities $\{\mu_f^{x0,(0)}\}_{x\in\mathcal{X}}$ and $\{\mu_m^{0y,(0)}\}_{y\in\mathcal{Y}}$. When the female and male population masses are normalized to one, f_x and m_y coincide with the corresponding population fractions.

Step 2: conditional choice probabilities and Pareto weights. For a female type $x=(e,r)$, define the partner-side inclusive value

$$\widetilde{S}_f^x(\lambda,m) = \exp\left(EV_f^{x0}/\sigma_\nu\right) + \sum_{y\in\mathcal{Y}} m_y \exp\left(\left[EV_f^{xy}(\lambda^{xy}) + \xi^{xy}\right]/\sigma_\nu\right). \tag{D.1}$$

Conditional on education e, the partner-choice probabilities are

$$q_f(0 \mid x;\lambda,m) = \frac{\exp\left(EV_f^{x0}/\sigma_v\right)}{\widetilde{S}_f^x(\lambda,m)}, \tag{D.2}$$

$$q_f(y \mid x;\lambda,m) = \frac{m_y \exp\left(\left[EV_f^{xy}(\lambda^{xy}) + \xi^{xy}\right]/\sigma_v\right)}{\widetilde{S}_f^x(\lambda,m)}, \qquad y \in \mathcal{Y}. \tag{D.3}$$

The education cost is common to all partner alternatives conditional on e and therefore cancels from these probabilities. Conditional on the education-cost vector v, the probability that a woman of fertility type r chooses education e is

$$p_f(e \mid r, v;\lambda,m) = \frac{\widetilde{S}_f^{(e,r)}(\lambda,m)\exp\left(-[c_{f,e}+v_e]/\sigma_v\right)}{\sum_{e'\in\mathcal{E}} \widetilde{S}_f^{(e',r)}(\lambda,m)\exp\left(-[c_{f,e'}+v_{e'}]/\sigma_v\right)}. \tag{D.4}$$

Integrating over the education-cost shocks gives the marginal education probability

$$P_f(x;\lambda,m) = \int p_f(e \mid r, v;\lambda,m) \prod_{e'\in\mathcal{E}} \phi(v_{e'}/\sigma_{ec}) \frac{dv_{e'}}{\sigma_{ec}}, \qquad x = (e,r), \tag{D.5}$$

where ϕ is the standard normal density. The integral is evaluated by Gauss–Hermite quadrature. The joint choice probabilities therefore factor as

$$\mu_f^{xy}(\lambda) = P_f(x;\lambda,m) q_f(y \mid x;\lambda,m), \qquad y \in \{0\} \cup \mathcal{Y}. \tag{D.6}$$

For men, define analogously

$$\widetilde{S}_m^y(\lambda,f) = \exp\left(EV_m^{0y}/\sigma_v\right) + \sum_{x\in\mathcal{X}} f_x \exp\left(\left[EV_m^{xy}(\lambda^{xy}) + \xi^{xy}\right]/\sigma_v\right). \tag{D.7}$$

This yields the conditional partner probabilities

$$q_m(0 \mid y;\lambda,f) = \frac{\exp\left(EV_m^{0y}/\sigma_v\right)}{\widetilde{S}_m^y(\lambda,f)}, \tag{D.8}$$

$$q_m(x \mid y;\lambda,f) = \frac{f_x \exp\left(\left[EV_m^{xy}(\lambda^{xy}) + \xi^{xy}\right]/\sigma_v\right)}{\widetilde{S}_m^y(\lambda,f)}, \qquad x \in \mathcal{X}, \tag{D.9}$$

and a marginal education probability $P_m(y;\lambda,f)$ obtained by integrating the male counterpart of Equation (D.4) over the male education-cost shocks. Thus,

$$\mu_m^{xy}(\lambda) = P_m(y;\lambda,f) q_m(x \mid y;\lambda,f), \qquad x \in \{0\} \cup \mathcal{X}. \tag{D.10}$$

For fixed female type x, the education marginal cancels between the probability of match-

ing with a type-y man and the probability of remaining single. Consequently,

$$\sigma_\nu \left[\log \mu_f^{xy}(\lambda) - \log \mu_f^{x0}(\lambda)\right] = \sigma_\nu \log m_y + EV_f^{xy}(\lambda^{xy}) + \xi^{xy} - EV_f^{x0}. \tag{D.11}$$

The analogous quasi-supply equation for men is

$$\sigma_\nu \left[\log \mu_m^{xy}(\lambda) - \log \mu_m^{0y}(\lambda)\right] = \sigma_\nu \log f_x + EV_m^{xy}(\lambda^{xy}) + \xi^{xy} - EV_m^{0y}. \tag{D.12}$$

At iteration k, combine Equations (D.11) and (D.12) with the market-clearing condition $\mu_\mathcal{X} \pi_r \mu_f^{xy} = \mu_\mathcal{Y} \mu_m^{xy}$. For every $(x, y) \in \mathcal{X} \times \mathcal{Y}$, the updated Pareto weight $\lambda^{xy,(k+1)}$ solves

$$\begin{aligned} 0 = \mathcal{G}_k^{xy}(\lambda) := &\left[EV_f^{xy}(\lambda) - EV_f^{x0}\right] - \left[EV_m^{xy}(\lambda) - EV_m^{0y}\right] \\ &- \sigma_\nu \log\left(\frac{\mu_\mathcal{Y} \mu_m^{0y,(k)} f_x^{(k)}}{\mu_\mathcal{X} \pi_r \mu_f^{x0,(k)} m_y^{(k)}}\right). \end{aligned} \tag{D.13}$$

The common match value ξ^{xy} cancels from this equation. Under condition (A1) of Appendix E, $\mathcal{G}_k^{xy}$ is continuous and strictly increasing, with $\mathcal{G}_k^{xy}(\lambda) \to -\infty$ as $\lambda \to 0^+$ and $\mathcal{G}_k^{xy}(\lambda) \to +\infty$ as $\lambda \to 1^-$. Equation (D.13) therefore has a unique solution in $(0, 1)$ for each couple type, conditional on the current masses and single-choice probabilities.

Step 3: update single-choice probabilities and type masses. Using $\lambda^{(k+1)}$ and the iteration-k masses, update the single-choice probabilities as

$$\mu_f^{x0,(k+1)} = P_f(x; \lambda^{(k+1)}, m^{(k)}) q_f(0 \mid x; \lambda^{(k+1)}, m^{(k)}), \tag{D.14}$$

$$\mu_m^{0y,(k+1)} = P_m(y; \lambda^{(k+1)}, f^{(k)}) q_m(0 \mid y; \lambda^{(k+1)}, f^{(k)}). \tag{D.15}$$

Update the education-type masses using the same education marginals:

$$f_{(e,r)}^{(k+1)} = \mu_\mathcal{X} \pi_r P_f((e, r); \lambda^{(k+1)}, m^{(k)}), \tag{D.16}$$

$$m_y^{(k+1)} = \mu_\mathcal{Y} P_m(y; \lambda^{(k+1)}, f^{(k)}). \tag{D.17}$$

Return to Step 2 and repeat until the vectors of single-choice probabilities and education-type masses have converged. Conditional on convergence, Equation (D.13) implies that every marriage submarket clears, while the updates in Step 3 imply that the type masses and single-choice probabilities are consistent with the choices they induce. The limiting objects therefore satisfy the equilibrium conditions in Definition 1.

E Proof of Equilibrium Existence

Following Gayle and Shephard (2019), we state sufficient conditions on the optimized values of marriage under which an equilibrium exists.[29] The proof in Gayle and Shephard (2019)

[29] The parametric household specification in Section 3 satisfies these conditions: nested-logit aggregation makes the spouse values continuously differentiable on the interior, Pareto efficiency gives the spouse-value derivative

establishes existence of Pareto weights that clear the marriage market for fixed type masses. Because education is endogenous in our model, we extend their fixed-point argument to include the female and male education-type masses.

Throughout, let $\mathcal{X} \equiv \mathcal{E} \times \mathcal{R}$ and $\mathcal{Y} \equiv \mathcal{E}$ denote the finite sets of female and male types. For $x = (e, r) \in \mathcal{X}$, r denotes the woman's fertility type. The following conditions are sufficient for the existence result.

(A1) For every $(x, y) \in \mathcal{X} \times \mathcal{Y}$, the spouse values $EV_f^{xy}, EV_m^{xy} : (0,1) \to \mathbb{R}$ are continuously differentiable and satisfy

$$\frac{\partial EV_f^{xy}(\lambda^{xy})}{\partial \lambda^{xy}} > 0, \qquad \frac{\partial EV_m^{xy}(\lambda^{xy})}{\partial \lambda^{xy}} < 0.$$

The endpoint values $EV_f^{xy}(1^-)$ and $EV_m^{xy}(0^+)$ are finite, whereas

$$\lim_{\lambda^{xy} \to 0^+} EV_f^{xy}(\lambda^{xy}) = -\infty, \qquad \lim_{\lambda^{xy} \to 1^-} EV_m^{xy}(\lambda^{xy}) = -\infty.$$

(A2) The population masses satisfy $\mu_{\mathcal{X}}, \mu_{\mathcal{Y}} > 0$, and the fertility-type shares satisfy $\pi_r > 0$ and $\sum_{r \in \mathcal{R}} \pi_r = 1$. The single values EV_f^{x0} and EV_m^{0y}, the deterministic match values ξ^{xy}, and the deterministic education costs are finite, and singlehood is always available.

(A3) Conditional on observable types, the matching shocks are independent Type-I extreme-value random variables with scale $\sigma_\nu > 0$ and are independent of the education-cost shocks. Every component of the education-cost vector is finite almost surely.

Choice probabilities and excess demand. For notational convenience, define

$$W_f^{xy}(\lambda) = \begin{cases} EV_f^{xy}(\lambda^{xy}) + \xi^{xy}, & y \in \mathcal{Y}, \\ EV_f^{x0}, & y = 0, \end{cases} \qquad W_m^{xy}(\lambda) = \begin{cases} EV_m^{xy}(\lambda^{xy}) + \xi^{xy}, & x \in \mathcal{X}, \\ EV_m^{0y}, & x = 0. \end{cases}$$

Conditional on own type, education costs cancel from the partner choice. After integrating over the education-cost shocks, the joint choice probabilities in Equations (28) and (29) therefore factor as

$$\mu_f^{xy} = P_f(x; \lambda, m) q_f(y \mid x; \lambda, m), \qquad \mu_m^{xy} = P_m(y; \lambda, f) q_m(x \mid y; \lambda, f),$$

signs, and log private consumption together with $s_{ij}(\lambda^{xy}) = \lambda^{xy}$ gives the stated endpoint behavior.

where P_f and P_m are the education marginals obtained by integrating over the education-cost shocks and

$$q_f(y \mid x; \lambda, m) = \frac{m_y \exp\left(W_f^{xy}(\lambda)/\sigma_\nu\right)}{\sum_{y' \in \{0\} \cup \mathcal{Y}} m_{y'} \exp\left(W_f^{xy'}(\lambda)/\sigma_\nu\right)},$$

$$q_m(x \mid y; \lambda, f) = \frac{f_x \exp\left(W_m^{xy}(\lambda)/\sigma_\nu\right)}{\sum_{x' \in \{0\} \cup \mathcal{X}} f_{x'} \exp\left(W_m^{x'y}(\lambda)/\sigma_\nu\right)}.$$

Here and below, $m_0 = f_0 = 1$ is the normalization for the single alternative. The type masses induced by the education choices are

$$f_x = \mu_{\mathcal{X}} \pi_r P_f(x; \lambda, m), \qquad m_y = \mu_{\mathcal{Y}} P_m(y; \lambda, f).$$

For $(x, y) \in \mathcal{X} \times \mathcal{Y}$, define excess demand by

$$ED^{xy}(\lambda, f, m) := m_y q_m(x \mid y; \lambda, f) - f_x q_f(y \mid x; \lambda, m). \tag{E.1}$$

At masses consistent with the education choices, $ED^{xy} = 0$ is equivalent to the market-clearing condition $\mu_{\mathcal{X}} \pi_r \mu_f^{xy} = \mu_{\mathcal{Y}} \mu_m^{xy}$.

Step 1: compact sets of positive type masses. First consider the full mass simplices

$$\overline{\mathcal{F}} = \left\{ f \geq 0 : \sum_{e \in \mathcal{E}} f_{(e,r)} = \mu_{\mathcal{X}} \pi_r \text{ for every } r \in \mathcal{R} \right\}, \qquad \overline{\mathcal{M}} = \left\{ m \geq 0 : \sum_{y \in \mathcal{Y}} m_y = \mu_{\mathcal{Y}} \right\}.$$

By (A1) and (A2), each exponentiated married value in the conditional logit formulas has a continuous extension to the closed Pareto-weight cube: it equals zero at the endpoint where the corresponding spouse value tends to minus infinity and has a finite limit at the opposite endpoint. Because the single alternative has finite value and unit mass, each education's partner inclusive value is finite and continuous as (λ, f, m) varies over $[0,1]^{\mathcal{X} \times \mathcal{Y}} \times \overline{\mathcal{F}} \times \overline{\mathcal{M}}$. This domain is compact, so the systematic education values range over a compact set. For every finite realization of the education-cost vector, the conditional logit probability of each education is continuous and strictly positive because that education's single alternative is always available. Integration preserves continuity by dominated convergence. Compactness therefore implies that the probability of every education is bounded below by a common constant $\phi_0 > 0$. Taking ϕ_0 smaller if necessary, define

$$\mathcal{F} = \left\{ f \in \overline{\mathcal{F}} : f_{(e,r)} \geq \mu_{\mathcal{X}} \pi_r \phi_0 \text{ for every } (e, r) \in \mathcal{X} \right\},$$
$$\mathcal{M} = \left\{ m \in \overline{\mathcal{M}} : m_y \geq \mu_{\mathcal{Y}} \phi_0 \text{ for every } y \in \mathcal{Y} \right\}.$$

These sets are nonempty, compact, and convex, and the education-mass update maps into $\mathcal{F} \times \mathcal{M}$.

Step 2: excess-demand derivatives and boundary signs. Fix $(f, m) \in \mathcal{F} \times \mathcal{M}$. By (A1), increasing λ^{xy} raises the woman's value of match (x, y) and lowers the man's value of that match. The logit formulas then imply

$$\frac{\partial ED^{xy}}{\partial \lambda^{xy}} < 0, \tag{E.2}$$

$$\frac{\partial ED^{xy'}}{\partial \lambda^{xy}} > 0 \qquad \text{for } y' \neq y, \tag{E.3}$$

$$\frac{\partial ED^{x'y}}{\partial \lambda^{xy}} > 0 \qquad \text{for } x' \neq x, \tag{E.4}$$

$$\frac{\partial ED^{x'y'}}{\partial \lambda^{xy}} = 0 \qquad \text{for } x' \neq x,\ y' \neq y. \tag{E.5}$$

This is the gross-substitutes structure used by Gayle and Shephard (2019).

As $\lambda^{xy} \to 0^+$, (A1) implies that the exponentiated female value for match (x, y) tends to zero, whereas the exponentiated male value tends to a finite, strictly positive limit. On $\mathcal{F} \times \mathcal{M}$, the relevant numerator masses are bounded away from zero, the logit denominators are bounded above, and their single-alternative terms are bounded away from zero. It follows that $f_x q_f(y \mid x) \to 0$ uniformly while $m_y q_m(x \mid y)$ is uniformly bounded below by a strictly positive constant, and hence $ED^{xy} > 0$ at the lower boundary. The symmetric argument gives $ED^{xy} < 0$ as $\lambda^{xy} \to 1^-$. Both bounds are uniform in the remaining Pareto weights and in $(f, m) \in \mathcal{F} \times \mathcal{M}$. Consequently, there is a common $\delta \in (0, 1/2)$ such that

$$ED^{xy}(\lambda, f, m) > 0 \quad \text{if } \lambda^{xy} = \delta, \qquad ED^{xy}(\lambda, f, m) < 0 \quad \text{if } \lambda^{xy} = 1 - \delta \tag{E.6}$$

for every (x, y), every value of the remaining coordinates in $[\delta, 1 - \delta]$, and every $(f, m) \in \mathcal{F} \times \mathcal{M}$. Let

$$C_\delta = [\delta, 1 - \delta]^{\mathcal{X} \times \mathcal{Y}}.$$

On the compact set $C_\delta \times \mathcal{F} \times \mathcal{M}$, define

$$L = \max_{(x,y) \in \mathcal{X} \times \mathcal{Y}} \sup_{C_\delta \times \mathcal{F} \times \mathcal{M}} \left| \frac{\partial ED^{xy}}{\partial \lambda^{xy}} \right|.$$

Continuity and (E.2) imply $0 < L < \infty$. Choose $\tau \in (0, L^{-1}]$.

Step 3: a joint fixed-point map. Define

$$T : C_\delta \times \mathcal{F} \times \mathcal{M} \longrightarrow C_\delta \times \mathcal{F} \times \mathcal{M}$$

by $T(\lambda, f, m) = (\widetilde{\lambda}, \widetilde{f}, \widetilde{m})$, where

$$\begin{aligned} \widetilde{\lambda}^{xy} &= \lambda^{xy} + \tau ED^{xy}(\lambda, f, m), \\ \widetilde{f}_{(e,r)} &= \mu_{\mathcal{X}} \pi_r P_f((e, r); \lambda, m), \\ \widetilde{m}_y &= \mu_{\mathcal{Y}} P_m(y; \lambda, f). \end{aligned}$$

The map is continuous. Continuity of the Pareto-weight component follows from (A1) and the logit formulas, while continuity of the education-mass components follows from dominated convergence.

It remains to verify that T is a self-map. For fixed (f, m), equations (E.2)–(E.5) and the choice of τ imply that every partial derivative of $\widetilde{\lambda}$ with respect to λ is nonnegative. Thus $\lambda \mapsto \widetilde{\lambda}$ is isotone on C_δ. Writing $\mathbf{1}$ for the all-ones vector, isotonicity and (E.6) yield, for every $\lambda \in C_\delta$,

$$\delta < \delta + \tau ED^{xy}(\delta\mathbf{1}, f, m) \leq \widetilde{\lambda}^{xy} \leq (1-\delta) + \tau ED^{xy}((1-\delta)\mathbf{1}, f, m) < 1 - \delta.$$

Therefore $\widetilde{\lambda} \in C_\delta$. Step 1 implies $\widetilde{f} \in \mathcal{F}$ and $\widetilde{m} \in \mathcal{M}$. Hence T is a continuous self-map of a nonempty compact convex set.

Step 4: existence. By Brouwer's fixed-point theorem, T has a fixed point (λ^*, f^*, m^*). The fixed-point condition for the Pareto weights implies $ED(\lambda^*, f^*, m^*) = 0$, so every marriage submarket clears. The fixed-point conditions for f^* and m^* make the education-type masses equal to the marginals generated by the joint education and partner choices. The choice probabilities in equations (28) and (29) are generated by optimal joint education-partner choices at (λ^*, f^*, m^*), and the allocations embedded in EV_f^{xy} and EV_m^{xy} solve the household problems. Thus all four conditions in Definition 1 hold. Finally, $\lambda^* \in C_\delta \subset (0,1)^{\mathcal{X}\times\mathcal{Y}}$, so every equilibrium Pareto weight is interior. □

F Proof of Identification

F.1 Proof of Proposition 2: Pareto Weights

Rearranging the quasi-demand and quasi-supply equations (D.11) and (D.12) derived in Appendix D, we obtain

$$\log \mu_f^{xy}(\lambda) - \log \mu_f^{x0}(\lambda) - \log m_y(\lambda) = \left[EV_f^{xy}(\lambda^{xy}) + \xi^{xy} - EV_f^{x0}\right]/\sigma_\nu, \tag{F.1}$$

$$\log \mu_m^{xy}(\lambda) - \log \mu_m^{0y}(\lambda) - \log f_x(\lambda) = \left[EV_m^{xy}(\lambda^{xy}) + \xi^{xy} - EV_m^{0y}\right]/\sigma_\nu. \tag{F.2}$$

The match and singlehood frequencies on the left-hand side are observed for men and recovered for women's latent fertility types using the first-stage posterior probabilities. Thus both left-hand sides are identified under the first-stage assumptions.

To identify the Pareto weights λ^{xy}, we require a *distribution factor* (Bourguignon et al., 2009): a variable that shifts the value of marriage only through the Pareto weight.[30] We use the market gender ratio $gr = \mu_\mathcal{X}/\mu_\mathcal{Y}$, a standard choice in the literature. A marginal perturbation in the

[30] If the education distribution is exogenous, the Pareto weights are identified from variation in the education distribution across markets (Gayle and Shephard, 2019).

gender ratio yields

$$\frac{\partial\left[\log\mu_f^{xy}(\lambda)-\log\mu_f^{x0}(\lambda)-\log m_y(\lambda)\right]}{\partial gr}=\frac{1}{\sigma_\nu}\frac{\partial EV_f^{xy}(\lambda^{xy})}{\partial\lambda^{xy}}\frac{\partial\lambda^{xy}}{\partial gr}, \tag{F.3}$$

$$\frac{\partial\left[\log\mu_m^{xy}(\lambda)-\log\mu_m^{0y}(\lambda)-\log f_x(\lambda)\right]}{\partial gr}=\frac{1}{\sigma_\nu}\frac{\partial EV_m^{xy}(\lambda^{xy})}{\partial\lambda^{xy}}\frac{\partial\lambda^{xy}}{\partial gr}, \tag{F.4}$$

where the gender ratio enters only through the Pareto weight λ^{xy}, leaving the single values EV_f^{x0} and EV_m^{0y} and the match value ξ^{xy} unaffected. This argument requires preferences, shock scales, and match values to be common across markets, though wages, divorce risk, and policy schedules may differ. The perturbation in gr is therefore taken holding the observed market characteristics fixed, and identification requires sex-ratio variation across markets conditional on the observed environment.

Provided $\partial\lambda^{xy}/\partial gr\neq 0$, and using $\partial EV_f^{xy}(\lambda^{xy})/\partial\lambda^{xy}>0$ and $\partial EV_m^{xy}(\lambda^{xy})/\partial\lambda^{xy}<0$, we take the ratio of Equations (F.3) and (F.4):

$$\begin{aligned}&\left(\frac{\partial\left[\log\mu_f^{xy}(\lambda)-\log\mu_f^{x0}(\lambda)-\log m_y(\lambda)\right]}{\partial gr}\right)^{-1}\frac{\partial\left[\log\mu_m^{xy}(\lambda)-\log\mu_m^{0y}(\lambda)-\log f_x(\lambda)\right]}{\partial gr}\\&\quad=\left(\frac{\partial EV_f^{xy}(\lambda^{xy})}{\partial\lambda^{xy}}\right)^{-1}\frac{\partial EV_m^{xy}(\lambda^{xy})}{\partial\lambda^{xy}}.\end{aligned} \tag{F.5}$$

All objects on the left-hand side are observed in the data, so

$$\pi^{xy}\equiv\left(\frac{\partial EV_f^{xy}(\lambda^{xy})}{\partial\lambda^{xy}}\right)^{-1}\frac{\partial EV_m^{xy}(\lambda^{xy})}{\partial\lambda^{xy}}$$

is identified.

A second restriction follows from efficient allocation within the couple. Let $EV^{xy}(\lambda^{xy})$ denote the couple's expected value, that is, the inclusive value $\sigma_\varepsilon c+\sigma_\varepsilon\ln\big(\sum_{n'=0}^{3}S_{n'}^{\rho_\varepsilon}(\lambda^{xy})\big)$ in Equation (18):

$$EV^{xy}(\lambda^{xy})\equiv\mathbb{E}\Big[\max_{a\in\mathcal{A}^M}\left\{V^{xy}(a;\lambda^{xy})+\varepsilon_{ij}(a)\right\}\Big]. \tag{F.6}$$

Substituting the household value function $V^{xy}(a;\lambda^{xy})=\lambda^{xy}V_f^{xy}(a;\lambda^{xy})+(1-\lambda^{xy})V_m^{xy}(a;\lambda^{xy})$ from Equation (16), and using that the shock $\varepsilon_{ij}(a)$ is shared across spouses, the couple's expected value aggregates the spouses' expected values:

$$EV^{xy}(\lambda^{xy})=\lambda^{xy}\,EV_f^{xy}(\lambda^{xy})+(1-\lambda^{xy})\,EV_m^{xy}(\lambda^{xy}). \tag{F.7}$$

The Pareto weight enters the couple's program only through the objective: the consumption share $s_{ij}(\lambda^{xy})$ in Equation (14) is itself the maximizer of the same weighted utility (13), so by the envelope theorem its indirect effect cancels, while the choice set $\mathcal{A}^M$, the budget constraint (10), and the divorce continuation values are all independent of λ^{xy}. The envelope property

of the nested-logit inclusive value therefore leaves only the explicit weight on each spouse's value:

$$\frac{\partial EV^{xy}(\lambda^{xy})}{\partial \lambda^{xy}} = \sum_a \mathbb{P}^{xy}(a;\lambda^{xy}) \left[V_f^{xy}(a;\lambda^{xy}) - V_m^{xy}(a;\lambda^{xy})\right] = EV_f^{xy}(\lambda^{xy}) - EV_m^{xy}(\lambda^{xy}), \quad \text{(F.8)}$$

where the last equality uses the cancellation of the shared shock in the difference. Differentiating (F.7) directly and equating with (F.8), the level terms $EV_f^{xy} - EV_m^{xy}$ cancel and we obtain

$$\lambda^{xy} \frac{\partial EV_f^{xy}(\lambda^{xy})}{\partial \lambda^{xy}} + (1-\lambda^{xy}) \frac{\partial EV_m^{xy}(\lambda^{xy})}{\partial \lambda^{xy}} = 0. \quad \text{(F.9)}$$

Equation (F.9) is the tangency condition between the iso-welfare line with weights $(\lambda^{xy}, 1-\lambda^{xy})$ and the couple's utility-possibility frontier. Combining (F.9) with the definition of π^{xy} yields

$$\lambda^{xy} = \frac{\pi^{xy}}{\pi^{xy}-1}, \quad \text{(F.10)}$$

which establishes identification of λ^{xy}.

F.2 Proof of Proposition 3: Preference Parameters

Identification is semiparametric: we do not impose the functional forms of (1) and (8)–(9), and we recover the per-gender, per-couple-type non-consumption utilities as objects, although we provide the specific utility function in the main text. Throughout, the Pareto weights $\{\lambda^{xy,k}\}$ are taken as identified (Subsection F.1); the superscript k indexes markets, suppressed when a single market suffices.

The architecture follows Gayle and Shephard (2019): almost all primitives are identified within a single market, and only the separation of the spouses' child preferences requires across-market variation. We impose the following assumptions:

Assumption ID-1: *The single-state shocks $\varepsilon_i(\cdot), \varepsilon_j(\cdot)$ are Type-I extreme value with location zero and scale $\sigma_\varepsilon > 0$. The couple shock $\varepsilon_{ij}(\cdot)$ is nested logit with scale σ_ε and nest-correlation $\rho_\varepsilon \in (0,1]$, the nests indexed by the number of children n, and is shared across spouses. These shock families are mutually independent.*

Assumption ID-2: *Flow utility is additively separable into a private-consumption component and a non-consumption component depending on labor supply and the number of children, and the consumption sub-utility $u_q(\cdot)$ is common across gender and marital status:*

$$u_f^{x0}(a) = u_q(q_f) + u_{f,n}^{x0}(l,n), \qquad u_m^{0y}(a) = u_q(q_m) + u_{m,n}^{0y}(l,n),$$
$$u_f^{xy}(a) = u_q(q_f) + u_{f,n}^{xy}(l,n), \qquad u_m^{xy}(a) = u_q(q_m) + u_{m,n}^{xy}(l,n).$$

The non-consumption utility is normalized to vanish at non-participation without children, $u_{f,0}^{x0}(0) = u_{m,0}^{0y}(0) = 0$; this is a free location normalization, since labor supply enters the choice probabilities only

through differences.

Assumption ID-3: *There is a known consumption level $\hat{q}$ with $\partial u_q(\hat{q})/\partial q = 1$.*

Assumption ID-4: *u_q is continuously differentiable, increasing, and concave in q.*

Assumption ID-5: *Child sub-utilities do not depend on the partner's type, current or former. No children are born outside marriage, so a never-married single is childless, and a divorced parent retains the children of the dissolved marriage.*

Assumption ID-6: *The divorce probabilities $d^{xy}(n) \in [0,1)$ and the budget sets are known. A childless agent's labor disutility does not depend on the wage, and the work-family conflict, which may depend on education, is absent without children.*

Assumption ID-7: *Childless single agents of each gender are observed across a wage support that includes a working bundle with consumption $\hat{q}$ from ID-3. Each couple type is observed at the base $(0,0,n)$, the joint-work profiles (l_i, l_j, n) with $l_i, l_j > 0$, and the marginal profiles $(l_i, 0, n), (0, l_j, n)$. For at least one observed joint-work cell, the mixed difference of the couple's consumption utility ϕ^{xy}, defined in Step 2, is nonzero. Each couple type xy is observed in at least two markets with distinct Pareto weights $\lambda^{xy,k} \neq \lambda^{xy,k'}$. Divorced parents of each gender are observed at every child count $n > 0$ choosing each labor alternative. For each female type x and child count $n > 0$, at least two partner types $y \neq y'$ have distinct divorce probabilities, $d^{xy}(n) \neq d^{xy'}(n)$.*

Throughout, we write $\Gamma_{g,r}(n) \equiv u^{xy}_{g,n}(0,n)$ for the married child sub-utility, the non-consumption utility at own non-participation, and $\Gamma^s_{f,r}(n) \equiv u^{x0}_{f,n}(0,n)$ and $\Gamma^s_m(n) \equiv u^{0y}_{m,n}(0,n)$ for its divorced counterparts, with per-child increments $\gamma_{g,r}(k) \equiv \Gamma_{g,r}(k) - \Gamma_{g,r}(k-1)$. Suppressing the partner type in this notation reflects ID-5; suppressing the woman's own education follows the parameterization of the main text and is not used in the argument, which is conducted at a fixed female type $x = (e,r)$. The location normalization in ID-2 implies $\Gamma^s_{f,r}(0) = \Gamma^s_m(0) = 0$, and we normalize the married schedules likewise, $\Gamma_{g,r}(0) = 0$.

F.2.1 Step 1: Singles

For a single individual the problem (3) is multinomial logit with scale σ_ε. We apply the wage-variation argument of Gayle and Shephard (2019) to *childless* single agents, for whom the only non-consumption component is the labor disutility; by ID-6 this does not depend on the wage, so the wage is an excluded shifter that moves consumption while preferences are held fixed. Taking $l = 0$ as the base alternative, for a childless single woman with wage w,

$$\log \frac{\mathbb{P}^{x0}_f(l; n{=}0, w)}{\mathbb{P}^{x0}_f(0; n{=}0, w)} = \frac{u_q\big(q_f(l;w)\big) - u_q\big(q_f(0;w)\big)}{\sigma_\varepsilon} + \delta^s_f(l), \tag{F.11}$$

where $\delta_f^s(l) \equiv [u_{f,0}^{x0}(l) - u_{f,0}^{x0}(0)]/\sigma_\varepsilon$ is the labor-disutility contrast, independent of the wage by ID-6, and $q_f(l;w)$ is the known single budget (2). Differentiating in the wage and using that out-of-work consumption does not respond to it, $\partial q_f(0;w)/\partial w = 0$ (the non-participant earns nothing), the base term and $\delta_f^s(l)$ drop out:

$$\frac{\partial}{\partial w} \log \frac{\mathbb{P}_f^{x0}(l; n{=}0, w)}{\mathbb{P}_f^{x0}(0; n{=}0, w)} = \frac{1}{\sigma_\varepsilon} u_q'\big(q_f(l;w)\big) \frac{\partial q_f(l;w)}{\partial w}, \tag{F.12}$$

with $\partial q_f(l;w)/\partial w$ known from the budget. The left-hand side is observed, so $g(q) \equiv u_q'(q)/\sigma_\varepsilon$ is identified at every working bundle traced as the wage varies (ID-7). ID-3 separates the scale: since $u_q'(\hat{q}) = 1$, evaluating at $\hat{q}$ gives the explicit forms

$$\sigma_\varepsilon = \frac{1}{g(\hat{q})}, \qquad u_q'(q) = \frac{g(q)}{g(\hat{q})}, \tag{F.13}$$

and integrating (F.13) identifies u_q over the consumption range spanned by the wage support, up to an immaterial additive location. With $(u_q, \sigma_\varepsilon)$ known, the level of (F.11) identifies the childless labor-disutility schedule $\delta_f^s(l)$, which the normalization $u_{f,0}^{x0}(0) = 0$ (ID-2) turns into the level $u_{f,0}^{x0}(l) = \sigma_\varepsilon \delta_f^s(l)$, so the childless single woman's value is identified as a level; the analogous argument for childless single men identifies the male labor disutility. The work-family conflict $\alpha_f(n;e)$ may depend on education, but it enters only the with-children states; it is recovered there given $(u_q, \sigma_\varepsilon)$ — for partnered women in Step 2, per couple type $x = (e, r)$, where its education-dependence is innocuous.

F.2.2 Step 2: Couples within a nest — ρ_ε and the non-consumption utilities

The wage-shifter argument of Step 1 does not extend to couples: moving the wife's education to shift her wage also moves the couple type and hence the Pareto weight λ^{xy}, and the same is true across markets, where the equilibrium weight $\lambda^{xy,k}$ varies with the wage. We therefore identify the couple primitives holding the couple type — and thus λ^{xy} — fixed, using the consumption variation generated by the spouses' labor choices. The consumption sub-utility u_q itself is not re-identified; it transfers from Step 1 under ID-2. Fix a couple type xy and a nest n. Conditional on the nest, the household problem (17) is logit with scale $\sigma_\varepsilon \rho_\varepsilon$, so for any labor profile relative to the base $(0,0)$,

$$\log \frac{\mathbb{P}^{xy}(l_i, l_j \mid n)}{\mathbb{P}^{xy}(0,0 \mid n)} = \frac{V^{xy}(l_i, l_j, n) - V^{xy}(0,0,n)}{\sigma_\varepsilon \rho_\varepsilon}. \tag{F.14}$$

By the value-function definition preceding (16), $V^{xy}(a) = (1 - d^{xy}(n))\, u^{xy}(a; \lambda^{xy}) + d^{xy}(n)\, \overline{EV}^{xy}(n)$ with $\overline{EV}^{xy}(n) \equiv \lambda^{xy} EV_f^{x0}(n) + (1 - \lambda^{xy}) EV_m^{0y}(n)$ depending on n alone, so the continuation term cancels from (F.14):

$$\log \frac{\mathbb{P}^{xy}(l_i, l_j \mid n)}{\mathbb{P}^{xy}(0,0 \mid n)} = \frac{(1 - d^{xy}(n))\big[u^{xy}(l_i, l_j, n) - u^{xy}(0,0,n)\big]}{\sigma_\varepsilon \rho_\varepsilon}. \tag{F.15}$$

Identification of ρ_ε. Fix one couple type xy in one market. With u_q, σ_ε known (Step 1) and λ^{xy}

known, the consumption component $\phi^{xy}(a) \equiv \lambda^{xy} u_q(\psi s_{ij} q(a)) + (1-\lambda^{xy}) u_q(\psi(1-s_{ij})q(a))$ is known at every alternative, where the share $s_{ij} = s_{ij}(\lambda^{xy}; q(a))$ solves the first-order condition of (13), yielding (14), at $q(a)$. By ID-2 the non-consumption utility is additively separable across spouses and the child sub-utility cancels within a nest, so (F.15) takes the form

$$L^{xy}(l_i, l_j, n) = \omega\left[\phi^{xy}(l_i, l_j, n) - \phi^{xy}(0, 0, n)\right] + a_{l_i} + b_{l_j}, \qquad \omega \equiv \frac{1 - d^{xy}(n)}{\sigma_\varepsilon \rho_\varepsilon}, \tag{F.16}$$

where $L^{xy}(l_i, l_j, n) \equiv \log[\mathbb{P}^{xy}(l_i, l_j \mid n)/\mathbb{P}^{xy}(0, 0 \mid n)]$ is observed, $a_{l_i} \equiv \omega\lambda^{xy}\delta_f(l_i, n)$, $b_{l_j} \equiv \omega(1-\lambda^{xy})\delta_m(l_j, n)$, $a_0 = b_0 = 0$, and $\delta_g(\cdot, n)$ denotes the spouse's labor and work–family contrast. For any joint-work cell (l_i, l_j) with $l_i, l_j > 0$, the spouse-separable terms drop out of the mixed difference, since a_{l_i} is common to $L^{xy}(l_i, l_j, n)$ and $L^{xy}(l_i, 0, n)$ while b_{l_j} is common to $L^{xy}(l_i, l_j, n)$ and $L^{xy}(0, l_j, n)$, and $L^{xy}(0, 0, n) = 0$ by construction, leaving

$$\omega = \frac{L^{xy}(l_i, l_j, n) - L^{xy}(l_i, 0, n) - L^{xy}(0, l_j, n)}{\phi^{xy}(l_i, l_j, n) - \phi^{xy}(l_i, 0, n) - \phi^{xy}(0, l_j, n) + \phi^{xy}(0, 0, n)}. \tag{F.17}$$

The numerator is observed and the denominator is known and, by ID-7, nonzero: pooled household income makes ϕ^{xy} a concave, non-additive function of the two earnings $(w_i l_i, w_j l_j)$, so the second difference does not vanish. Thus ω is identified from a single couple type in a single market, and inverting $\omega = (1 - d^{xy}(n))/(\sigma_\varepsilon \rho_\varepsilon)$ gives the explicit form

$$\rho_\varepsilon = \frac{1 - d^{xy}(n)}{\sigma_\varepsilon\, \omega}, \tag{F.18}$$

with σ_ε from (F.13) and $d^{xy}(n)$ known.

Non-consumption utilities. Given ρ_ε, the coefficient $\omega^{xy}(n) \equiv (1 - d^{xy}(n))/(\sigma_\varepsilon \rho_\varepsilon)$ is known for every couple type and nest, and the marginal cells in (F.16) give $a_{l_i} = L^{xy}(l_i, 0, n) - \omega[\phi^{xy}(l_i, 0, n) - \phi^{xy}(0, 0, n)]$ and $b_{l_j} = L^{xy}(0, l_j, n) - \omega[\phi^{xy}(0, l_j, n) - \phi^{xy}(0, 0, n)]$, hence

$$\delta_f(l_i, n) = \frac{a_{l_i}}{\omega\lambda^{xy}} = u^{xy}_{f,n}(l_i, n) - u^{xy}_{f,n}(0, n), \qquad \delta_m(l_j, n) = \frac{b_{l_j}}{\omega(1-\lambda^{xy})} = u^{xy}_{m,n}(l_j, n) - u^{xy}_{m,n}(0, n). \tag{F.19}$$

The labor-disutility and partnered work–family components of each spouse's non-consumption utility are therefore identified for every n from a single market by ID-7. What (F.19) leaves undetermined is the within-nest base level $u^{xy}_{f,n}(0, n)$ and $u^{xy}_{m,n}(0, n)$, which at $l = 0$, where labor disutility and work-family conflict vanish, equal the child sub-utilities $\Gamma_{f,r}(n)$ and $\Gamma_m(n)$; these are levels that difference out within a nest and are recovered across nests in Step 3.

F.2.3 Step 3: Couples across nests — child sub-utilities

The across-nest probabilities satisfy $\log[\mathbb{P}^{xy}(n)/\mathbb{P}^{xy}(0)] = (I^{xy}_n - I^{xy}_0)/\sigma_\varepsilon$, where $I^{xy}_n \equiv \sigma_\varepsilon \rho_\varepsilon \log S_n(\lambda^{xy})$ is the nest inclusive value in (18). Hence $I^{xy}_n - I^{xy}_0$ is identified. Writing the within-nest sum

relative to its base and using $\sum_{l_i,l_j} \mathbb{P}^{xy}(l_i, l_j \mid n) = 1$,

$$I_n^{xy} = d^{xy}(n)\,\overline{EV}^{xy}(n) + (1 - d^{xy}(n))\,u^{xy}(0,0,n) \;-\; \sigma_\varepsilon \rho_\varepsilon \log \mathbb{P}^{xy}(0,0 \mid n), \tag{F.20}$$

in which the last term is observed and, by Step 2, $u^{xy}(0,0,n) = \phi^{xy}(0,0,n) + \lambda^{xy}\Gamma_{f,r}(n) + (1 - \lambda^{xy})\Gamma_m(n)$ with $\phi^{xy}(0,0,n)$ known. Under ID-5, divorced parents retain the n children of the marriage, so their continuation values include the costs and utility of those children. For a single parent of either gender the value is the scale-σ_ε expected maximum over labor; Hotz–Miller inversion of the observed divorced-parent CCPs, given $(u_q, \sigma_\varepsilon)$, identifies it up to the additive divorced child sub-utility, which under ID-5 is a free schedule:

$$EV_f^{x0}(n) = V_f^{x0}(n) + \Gamma_{f,r}^s(n), \qquad EV_m^{0y}(n) = V_m^{0y}(n) + \Gamma_m^s(n),$$

with the non-child parts $V_f^{x0}(n), V_m^{0y}(n)$ known. Substituting the continuation into (F.20) and collecting the unknowns, the married sub-utilities enter the surviving-marriage flow with weight $1 - d^{xy}(n)$, while the divorced sub-utilities enter the continuation with weight $d^{xy}(n)$. Every other term — $(1 - d^{xy}(n))\,\phi^{xy}(0,0,n)$, the known continuation parts $d^{xy}(n)V_f^{x0}(n)$, $d^{xy}(n)V_m^{0y}(n)$, and the observed CCP term — is known; anchoring at $n = 0$, where all child sub-utilities vanish, yields the identified object

$$\Psi^{xy}(n) \;=\; \lambda^{xy}\,A^{xy}(n) + (1 - \lambda^{xy})\,B^{xy}(n), \qquad n > 0, \tag{F.21}$$

where $\Psi^{xy}(n)$ is the (identified) difference $I_n^{xy} - I_0^{xy}$ net of all known terms and

$$A^{xy}(n) \equiv (1 - d^{xy}(n))\,\Gamma_{f,r}(n) + d^{xy}(n)\,\Gamma_{f,r}^s(n), \qquad B^{xy}(n) \equiv (1 - d^{xy}(n))\,\Gamma_m(n) + d^{xy}(n)\,\Gamma_m^s(n) \tag{F.22}$$

are each spouse's child sub-utility averaged over the survival and dissolution of the marriage. At this stage $\Psi^{xy}(n)$ is identified, while the composites $A^{xy}(n)$ and $B^{xy}(n)$ are the unknowns to be recovered.

Equation (F.21) is one equation in the two composites $A^{xy}(n), B^{xy}(n)$, which enter only through the scalar λ^{xy}; separating them requires a second value of λ^{xy} at the same couple type. For fixed xy the equilibrium weight $\lambda^{xy,k}$ varies across markets k through the sex ratio (Subsection F.1), while the co-varying wages are known and net out of $\Psi^{xy,k}(n)$; divorce risk for a given couple type and number of children is common across markets, so the composites are market-invariant.[31] By (F.21), in each market k

$$\Psi^{xy,k}(n) = \lambda^{xy,k}\,A^{xy}(n) + (1 - \lambda^{xy,k})\,B^{xy}(n). \tag{F.23}$$

Two markets k, k' with $\lambda^{xy,k} \neq \lambda^{xy,k'}$ give a 2×2 system with determinant $\lambda^{xy,k} - \lambda^{xy,k'} \neq 0$,

[31] If divorce risk also varied across markets, the four schedules would still be identified: (F.21) is then linear in $(\Gamma_{f,r}, \Gamma_{f,r}^s, \Gamma_m, \Gamma_m^s)(n)$ with known coefficients $(\lambda(1-d), \lambda d, (1-\lambda)(1-d), (1-\lambda)d)$, and no linear functional annihilates this vector over $(\lambda, d) \in (0,1)^2$, so four market–partner-type combinations with a full-rank coefficient matrix suffice.

solved explicitly as

$$A^{xy}(n) = \frac{(1-\lambda^{xy,k'})\,\Psi^{xy,k}(n) - (1-\lambda^{xy,k})\,\Psi^{xy,k'}(n)}{\lambda^{xy,k} - \lambda^{xy,k'}}, \tag{F.24}$$

$$B^{xy}(n) = \frac{\lambda^{xy,k}\,\Psi^{xy,k'}(n) - \lambda^{xy,k'}\,\Psi^{xy,k}(n)}{\lambda^{xy,k} - \lambda^{xy,k'}}. \tag{F.25}$$

Every term on the right of (F.24)–(F.25) is identified, so the composites $A^{xy}(n)$ and $B^{xy}(n)$ are identified. The composites still mix the married and divorced schedules. By ID-5 none of the four schedules depends on the partner type, so $A^{xy}(n)$ and $B^{xy}(n)$ vary across y only through the divorce probability $d^{xy}(n)$. Two partner types $y \neq y'$ with $d^{xy}(n) \neq d^{xy'}(n)$ (ID-7) give a second 2×2 system, with determinant $d^{xy'}(n) - d^{xy}(n) \neq 0$:

$$\Gamma_{f,r}(n) = \frac{d^{xy'}(n)\,A^{xy}(n) - d^{xy}(n)\,A^{xy'}(n)}{d^{xy'}(n) - d^{xy}(n)}, \qquad \Gamma^{s}_{f,r}(n) = \frac{(1-d^{xy'}(n))\,A^{xy}(n) - (1-d^{xy}(n))\,A^{xy'}(n)}{d^{xy}(n) - d^{xy'}(n)}, \tag{F.26}$$

$$\Gamma_{m}(n) = \frac{d^{xy'}(n)\,B^{xy}(n) - d^{xy}(n)\,B^{xy'}(n)}{d^{xy'}(n) - d^{xy}(n)}, \qquad \Gamma^{s}_{m}(n) = \frac{(1-d^{xy'}(n))\,B^{xy}(n) - (1-d^{xy}(n))\,B^{xy'}(n)}{d^{xy}(n) - d^{xy'}(n)}. \tag{F.27}$$

Every quantity on the right is already identified: the composites from (F.24)–(F.25), the Pareto weights $\lambda^{xy,k}$ from Subsection F.1, and the divorce probabilities $d^{xy}(n)$ by ID-6; the four schedules $\Gamma_{f,r}, \Gamma^{s}_{f,r}, \Gamma_m, \Gamma^{s}_m$ are therefore identified. Additional markets and partner types over-identify the schedules.[32] Differencing gives the per-child marginal utilities $\gamma_{f,r}(k) = \Gamma_{f,r}(k) - \Gamma_{f,r}(k-1)$ and $\gamma_m(k) = \Gamma_m(k) - \Gamma_m(k-1)$. This establishes Proposition 3.[33] □

F.3 Proof of Proposition 4: Education Choice and Matching

Identification of σ_ν and ξ^{xy}. For each x, y, (F.1) and (F.2) give $A = (C + \xi^{xy})/\sigma_\nu$ and $B = (D + \xi^{xy})/\sigma_\nu$, where

$$A = \log \mu_f^{xy} - \log \mu_f^{x0} - \log m_y, \qquad C = EV_f^{xy}(\lambda^{xy}) - EV_f^{x0},$$
$$B = \log \mu_m^{xy} - \log \mu_m^{0y} - \log f_x, \qquad D = EV_m^{xy}(\lambda^{xy}) - EV_m^{0y}.$$

Subtracting the two equations removes the match-quality term:

$$A - B = \frac{C - D}{\sigma_\nu}. \tag{F.28}$$

[32] The parametric implementation restricts $\Gamma^{s}_{f,r} = \Gamma_{f,r}$, so a divorced mother values her children by the married schedule, and $\Gamma^{s}_m = 0$, so children do not enter a divorced father's utility. Both are admissible values of the identified schedules, so the restriction is testable rather than required for identification. Under it the composites reduce to $A^{xy}(n) = \Gamma_{f,r}(n)$ and $B^{xy}(n) = (1 - d^{xy}(n))\,\Gamma_m(n)$: divorce risk discounts the husband's child sub-utility but not the wife's, because the children remain with her.

[33] Under the parametric specification of the main text, the identified objects map one-to-one to the structural parameters: (F.19) delivers β_f, β_m and the partnered work–family schedules $\alpha_f(\cdot\,; e), \alpha_m(\cdot\,; y)$ of (8)–(9), while (F.26)–(F.27) deliver $\{\gamma_{f,r}(\cdot)\}_r$ and $\gamma_m(\cdot)$, which under the parameterization (11) pin γ_0, the linear coefficients $\{\gamma_{1,f,r}\}_r, \gamma_{1,m}$, and the common curvature γ_2 through $\gamma_{g,r}(k) = \gamma_0 \mathbf{1}\{k = 1\} + \gamma_{1,g,r} + \gamma_2(2k-1)$.

The left-hand side is identified from matching frequencies, and C, D are identified from household values and singlehood outcomes in the preceding steps. Provided at least one couple type has $A \neq B$ (equivalently, $C \neq D$),

$$\sigma_\nu = \frac{C - D}{A - B}, \qquad \xi^{xy} = \sigma_\nu A - C \tag{F.29}$$

identifies the matching-taste scale and then the match-quality term for every couple type, including cells for which $A = B$. The parameterization of ξ^{xy} is identified if its education terms, including any same-education indicator, have full rank across the observed couple types.

Identification of $c_{f,e}$ and the cost-shock distribution. At this stage $\{\lambda^{xy,k}\}$, $\{EV_g^{xy}(\lambda^{xy})\}$, $\{\xi^{xy}\}$, and σ_ν are identified, so for each situation (r,k) the cost-free education inclusive value

$$\tilde{S}_0^{(e,r),k} \equiv \sum_{y \in \{0\} \cup \mathcal{E}} m_y^k \exp\big([EV_f^{(e,r)y}(\lambda^{(e,r)y,k}) + \xi^{(e,r)y}]/\sigma_\nu\big) \tag{F.30}$$

is identified. Under the simultaneous game, the agent observes marriage taste shocks before choosing education and evaluates each education level at the realized maximum across partners. By the extreme-value max representation,

$$\max_{y \in \{0\} \cup \mathcal{E}} \big\{EV_f^{(e,r)y}(\lambda) + \xi^{(e,r)y} + \sigma_\nu \log m_y^k + \sigma_\nu \nu^{(e,r)y}\big\} = \sigma_\nu \log \tilde{S}_0^{(e,r),k} + \sigma_\nu c + \sigma_\nu \zeta_e^{(r,k)}, \tag{F.31}$$

where c is the Euler–Mascheroni constant, as in (18), and $\zeta_e^{(r,k)}$ is a mean-zero Type-I extreme value variate, the location-zero variate less its mean c. Because the match shock $\nu^{(e,r)y}$ varies with the pairing through $x = (e,r)$, distinct education levels carry independent shock vectors, so the inner-max fluctuations $\{\zeta_e^{(r,k)}\}_e$ are i.i.d. across e. The education choice is therefore

$$e^* = \arg\max_{e \in \mathcal{E}} \big\{\, z_e^{(r,k)} - c_{f,e} - \tilde{v}_e^{(r,k)} \,\big\}, \qquad z_e^{(r,k)} \equiv \sigma_\nu \log \tilde{S}_0^{(e,r),k}, \quad \tilde{v}_e^{(r,k)} \equiv v_e - \sigma_\nu \zeta_e^{(r,k)}, \tag{F.32}$$

the constant $\sigma_\nu c$ absorbed. The index $z_e^{(r,k)}$ is identified, since σ_ν, $\{\xi^{xy}\}$, $\{EV_g^{xy}\}$, and $\{m_y^k\}$ are.[34] Education choice is governed by $\tilde{v}_e^{(r,k)} = v_e - \sigma_\nu \zeta_e^{(r,k)}$. This shock combines the normal education-cost shock with the negative of a scaled mean-zero extreme-value shock; it is i.i.d. across e and has a positive density. The shifter $z_e^{(r,k)}$ varies across situations (r,k) through $EV_f^{(e,r)y}$ while $c_{f,e}$ does not; this variation is present within a single market across the fertility types r and is reinforced across markets. The observed education shares are $P_f^k((e,r)) = \Pr\big(z_e^{(r,k)} - c_{f,e} - \tilde{v}_e^{(r,k)} \geq z_{e'}^{(r,k)} - c_{f,e'} - \tilde{v}_{e'}^{(r,k)} \;\; \forall e'\big)$.

A single situation does not suffice. Within one situation the $|\mathcal{E}| - 1$ education shares confront $|\mathcal{E}|$ unknowns, the $|\mathcal{E}| - 1$ cost gaps and the scale σ_{ec}, one more than the moments; adding education levels adds shares and cost gaps in equal number and does not close the gap. What separates them is variation in the *known* shifter $z_e^{(r,k)}$ that holds the cost and the shock law fixed.

[34] The index is identical under the sequential reading of the game, in which education is chosen before the matching shocks are observed: the inclusive value $\sigma_\nu \log \tilde{S}_0^{(e,r),k}$ appears either way, and only the education-margin shock changes, from v_e to the composite $\tilde{v}_e^{(r,k)}$.

Variation across fertility types supplies an additional shift in the marriage market return to education. We recover type-specific education shares using the first-stage posterior probabilities. Thus, the cost levels and σ_{ec} are identified within a single market from the cross-r response of the education shares; across-market variation enlarges the set of z-gaps but is not required. In sum, identification requires the following assumption.

Assumption ID-8 (cross-situation variation). *The education costs and the distribution of the education-cost shocks are common across situations (r, k). For some pair of education levels $e \neq e'$ and some pair of situations $s \neq s'$, the shifter gaps differ: $z_e^s - z_{e'}^s \neq z_e^{s'} - z_{e'}^{s'}$. The resulting equation (F.34) has a unique positive solution for σ_{ec}.*

Equation (F.32) is a discrete choice whose mean utilities $z_e^{(r,k)} - c_{f,e}$ are known up to the costs and whose shocks $\tilde{v}_e^{(r,k)}$ are i.i.d. with a positive density. The share vector of a situation is a smooth function of the mean-utility gaps whose Jacobian has a positive diagonal and negative off-diagonals, and is therefore injective. For a situation s and a candidate scale $\sigma > 0$, let $M_{ee'}^s(\sigma)$ denote the mean-utility gaps implied by the observed education shares of situation s when the composite law is evaluated at σ; the composite is a known one-parameter family in σ because its extreme-value component carries the already-identified scale σ_v, so $M_{ee'}^s(\sigma)$ is a known function. At the true scale the inversion returns the true gaps:

$$M_{ee'}^s(\sigma_{ec}) = \left(z_e^s - z_{e'}^s\right) - \left(c_{f,e} - c_{f,e'}\right). \tag{F.33}$$

The cost gap is common across situations by ID-8, so differencing two situations $s \neq s'$ removes it: σ_{ec} solves the known scalar equation

$$M_{ee'}^s(\sigma) - M_{ee'}^{s'}(\sigma) = \left(z_e^s - z_{e'}^s\right) - \left(z_e^{s'} - z_{e'}^{s'}\right), \tag{F.34}$$

whose right-hand side is nonzero for some pair (e, e') and whose solution is unique by ID-8. Thus σ_{ec} is identified.[35] The costs then follow from (F.33) in a single line,

$$c_{f,e} = \left(z_e^s - z_1^s\right) - M_{e1}^s(\sigma_{ec}), \tag{F.35}$$

with the normalization $c_{f,1} = 0$,[36] so the education costs $\{c_{f,e}\}$ are identified. The analogous argument identifies $\{c_{m,y}\}$; for men, who carry no fertility type, the situations are markets alone, so the shifter variation must come across markets; alternatively, with the cost-shock scale common across genders as in the main text, the female-identified σ_{ec} transfers and a single market suffices. This establishes Proposition 4. □

[35] In estimation we restrict both the education costs and σ_{ec} to be common across the two cohorts. Cohorts then serve as additional situations in ID-8: cost gaps cancel between cohorts exactly as between situations within a cohort, so the across-cohort shifts in the marriage market returns, and hence in z, supply identifying variation for the common scale, with each cohort's composite law evaluated at its own identified σ_v. A single situation per cohort would already suffice, since two situations provide $2(|\mathcal{E}| - 1)$ share moments for the $|\mathcal{E}| - 1$ common cost gaps and the common scale. The within-cohort variation of ID-8 then becomes over-identifying, so the pooling restriction is testable.

[36] The estimates instead normalize costs relative to secondary education, an equivalent choice.

G Estimation Details

G.1 Recovering unobserved types

This subsection details how we recover the distribution of unobserved fertility types for women, following Arcidiacono et al. (2025). We allow for heterogeneity in women's intrinsic preferences over children through a binary latent *fertility type* $r \in \{L, H\}$. The first stage estimates a type share $\pi_{r,k,c}$ for each local marriage market k and cohort c, and a posterior type probability q_{ir} for each woman. Let k_i and c_i denote woman i's market and cohort. The type shares vary across market–cohort cells, while the remaining likelihood parameters θ_{type} are estimated jointly across cohorts. These first-stage objects are then treated as known inputs to the second-stage moment-based estimation, which recovers the remaining structural parameters in the spirit of Arcidiacono and Miller (2011). Subsection G.1.6 describes how the first-stage outputs feed into that second stage.

The first-stage auxiliary model has two ingredients. The first is a *measurement system*: variables that are informative about the latent fertility type and are used to identify its distribution. The second is a set of *flexible reduced-form equations* for the outcomes that enter the structural model. All remaining first-stage parameters are treated as nuisance parameters and discarded before the second stage begins; only $(q_{ir}, \pi_{r,k,c})$ are retained. The first stage uses the measurements and observed outcomes to estimate type shares and posterior type probabilities without imposing the structural choice model. The second stage conditions on these estimates.

G.1.1 Dependent variables used in the first stage

Two groups of dependent variables enter the first-stage likelihood.

The *measurements* are variables whose sole role in the first stage is to pin down the distribution of the unobserved fertility type. They do not enter the structural model directly. We use:

- age at first child,
- age at first relationship,
- number of siblings.

The economic rationale behind each measurement is discussed in Section 4. In short, conditional on observables, women with stronger preferences for children (high fertility type) tend to enter their first relationship and have their first child earlier in the life cycle, given that fertility rates are higher among couples than singles and that fecundity declines with age. The number of siblings captures intergenerational transmission of fertility preferences: women raised in larger families are more likely to be of high fertility type themselves. This measurement therefore provides an additional signal that does not rely on the woman's own fertility choices.

The *outcomes* are the endogenous variables that the structural model is designed to explain:

- education,

- marriage,
- labor supply,
- number of children.

Including these outcomes in the first stage ensures that the recovered type distribution is consistent with the same latent heterogeneity that governs behavior in the structural model, in line with the two-step approach introduced by Arcidiacono and Miller (2011).

G.1.2 Measurement system

Let b index a first-stage dependent variable. The measurement equation, conditional on fertility type $R_i = r$, is

$$y_{ib} = X_i\tilde{\beta}_b + \tilde{\gamma}_{br} + \epsilon_{ib}, \tag{G.1}$$

where X_i collects demographic and family-background covariates, parental education and parental income, while $\tilde{\gamma}_{br}$ is a measurement- and type-specific intercept capturing how the latent fertility type shifts the distribution of measurement b. As described in Section 2, outcomes are observed through age 38; the value $\bar{y}_b = 39$ codes non-occurrence of a timing event by that age. With two fertility types, we impose the normalization $\tilde{\gamma}_{b,L} = 0$ for every measurement.

Likelihood for timing measurements. The two age variables record *when* an event occurs, and are observed only if it occurs within the observation window; otherwise the variable is coded $y_{ib} = \bar{y}_b$. We therefore model each timing measurement conditional on the event occurring within the window, and write $d_{ib} = \mathbb{I}[y_{ib} < \bar{y}_b]$ for occurrence. Age at first child is observed exactly whenever it occurs; age at first relationship is also left-censored at $\underline{y}_b = 19$ (any age below 19 is coded as 19). Assuming $\epsilon_{ib} \sim \mathcal{N}(0, \sigma_b^2)$, the likelihood conditional on $R_i = r$ is

$$L^{ce}_{ibr}(\theta_{type}; y_{ib}, X_i) = \left[\left\{ \frac{1}{\sigma_b}\phi\left(\frac{y_{ib} - X_i\tilde{\beta}_b - \tilde{\gamma}_{br}}{\sigma_b} \right) \right\}^{\mathbb{I}[\underline{y}_b < y_{ib} < \bar{y}_b]} \times \left\{ \Phi\left(\frac{\underline{y}_b - X_i\tilde{\beta}_b - \tilde{\gamma}_{br}}{\sigma_b} \right) \right\}^{\mathbb{I}[y_{ib} = \underline{y}_b]} \right]^{d_{ib}}, \tag{G.2}$$

where $\phi(\cdot)$ and $\Phi(\cdot)$ denote the standard-normal density and CDF. A woman whose event does not occur within the window ($d_{ib} = 0$) contributes a factor of unity for that measurement. The left-censoring indicator is never activated for age at first child, so the same expression covers both variables.[37]

[37] The natural alternative is to treat $y_{ib} = \bar{y}_b$ as a right-censored observation contributing $1 - \Phi\big((\bar{y}_b - X_i\tilde{\beta}_b - \tilde{\gamma}_{br})/\sigma_b\big)$. We do not adopt this treatment because it would violate the conditional independence that the multiplicative first-stage likelihood (G.12) imposes and on which identification rests (Section 4). For a woman who never enters a relationship within the window, age at first relationship, age at first child, and the

Two conditions underlie the identification result invoked in Section 4. First, the measurements must be mutually independent conditional on the latent type and the covariates. This condition justifies the multiplicative form of $L_{ir}^{(m)}$, and it is the condition that the treatment above protects. Second, the measurements must be non-trivially informative about the type, in the sense that their type-conditional distributions differ.

Likelihood for ordered measurements. The number of siblings is modeled as an ordered categorical variable with four bins $\{0, 1, 2, 3+\}$ via an ordered logit specification. Indexing the bins by $l = 0, 1, \ldots, L_b$ with $L_b = 3$,

$$L_{ibr}^{o}\left(\theta_{type}; y_{ib}, X_i\right) = \prod_{l=0}^{L_b} P_{ibrl}^{\mathbb{I}[y_{ib}=l]}, \tag{G.3}$$

$$\begin{aligned} P_{ibrl} &= \Pr\left(\tilde{\kappa}_{l-1,b} < X_i\tilde{\beta}_b + \tilde{\gamma}_{br} + \epsilon_{ib} \leq \tilde{\kappa}_{l,b}\right) \\ &= \frac{1}{1+\exp\left(X_i\tilde{\beta}_b + \tilde{\gamma}_{br} - \tilde{\kappa}_{l,b}\right)} - \frac{1}{1+\exp\left(X_i\tilde{\beta}_b + \tilde{\gamma}_{br} - \tilde{\kappa}_{l-1,b}\right)}, \end{aligned} \tag{G.4}$$

with $\tilde{\kappa}_{-1,b} = -\infty$ and $\tilde{\kappa}_{L_b,b} = \infty$.

Measurement-system likelihood. Stacking across the three measurements, the joint measurement-system likelihood conditional on type r is

$$L_{ir}^{(m)} = \prod_{b\in\text{time}} L_{ibr}^{ce} \prod_{b\in\text{ord}} L_{ibr}^{o}, \tag{G.5}$$

where time denotes the two age variables, each contributing a factor only for the women whose event occurs within the window, and ord the number-of-siblings measurement.

G.1.3 Other first-stage equations

Estimation of the outcome equations proceeds analogously to the measurement system, with one addition. Following Arcidiacono et al. (2025), we include previous-decision dummies in the outcome equations to capture the sequential nature of the structural choices. For each outcome b and type r, the outcome equation is

$$y_{ib} = X_i\tilde{\beta}_{1b} + Z_{ib}\tilde{\beta}_{2b} + \tilde{\gamma}_{br} + \epsilon_{ib}, \tag{G.6}$$

where X_i denotes the same parental-background covariates as in the measurement system and Z_{ib} collects the outcome-specific previous-decision dummies. Specifically, Z_{ib} contains educa-

number of children are not three separate signals about her type: they are one event, recorded as 39, 39, and 0. Entering all three multiplicatively lets a single type-specific intercept account for the same fact three times, and a two-type mixture then attains a large likelihood gain by devoting one component entirely to this group, so that the recovered type would identify the extensive margin of family formation rather than the intensity of preferences it is meant to measure. Conditioning the timing measurements on occurrence lets this extensive margin inform the type exactly once, through the marriage equation and the zero bin of the number-of-children equation, whose supports overlap those of the rest of the sample. Equivalently, women whose events never occur are not dropped from the first stage: they remain in the likelihood and inform the type through the marriage and number-of-children equations, whereas dropping them would discard the extensive margin altogether and distort the recovered type shares.

tion dummies for the marriage equation, and education dummies together with a marriage indicator for the labor-supply and number-of-children equations; Z_{ib} is empty for the education equation itself. The coefficients $\tilde{\beta}_{2b}$ on the previous-decision dummies are outcome-specific but are *not* allowed to vary with the latent type r; only the intercept $\tilde{\gamma}_{br}$ carries type-specific heterogeneity. As before, we impose the normalization $\tilde{\gamma}_{b,L} = 0$.

Likelihood for binary outcomes. The marriage indicator equals one if the woman has a qualifying partnership under the definition in Section 2. Assuming ϵ_{ib} is consistent with a binary logit,

$$L_{ibr}^{bin}\left(\theta_{type}; y_{ib}, X_i, Z_{ib}\right) = P_{ibr}^{\mathbb{I}[y_{ib}=1]}\left(1 - P_{ibr}\right)^{\mathbb{I}[y_{ib}=0]}, \tag{G.7}$$

$$P_{ibr} = \frac{\exp\left(X_i\tilde{\beta}_{1b} + Z_{ib}\tilde{\beta}_{2b} + \tilde{\gamma}_{br}\right)}{1 + \exp\left(X_i\tilde{\beta}_{1b} + Z_{ib}\tilde{\beta}_{2b} + \tilde{\gamma}_{br}\right)}. \tag{G.8}$$

Likelihood for ordered outcomes. Education, labor supply, and number of children are all modeled as ordered categorical variables via ordered logit. For each such outcome b with $L_b + 1$ bins indexed by $l \in \{0, 1, \ldots, L_b\}$,

$$L_{ibr}^{o}\left(\theta_{type}; y_{ib}, X_i, Z_{ib}\right) = \prod_{l=0}^{L_b} P_{ibrl}^{\mathbb{I}[y_{ib}=l]}, \tag{G.9}$$

$$P_{ibrl} = \frac{1}{1 + \exp\left(X_i\tilde{\beta}_{1b} + Z_{ib}\tilde{\beta}_{2b} + \tilde{\gamma}_{br} - \tilde{\kappa}_{l,b}\right)} - \frac{1}{1 + \exp\left(X_i\tilde{\beta}_{1b} + Z_{ib}\tilde{\beta}_{2b} + \tilde{\gamma}_{br} - \tilde{\kappa}_{l-1,b}\right)}, \tag{G.10}$$

with $\tilde{\kappa}_{-1,b} = -\infty$ and $\tilde{\kappa}_{L_b,b} = \infty$. Education takes four bins ($L_b = 3$; primary, secondary, bachelor's, master's/PhD), labor supply takes three ($L_b = 2$; not working, part-time, full-time), and number of children takes four ($L_b = 3$; 0, 1, 2, 3+).

Likelihood for all outcome equations. Combining the binary and ordered components, the outcome-equation likelihood conditional on type r is

$$L_{ir}^{(o)} = \prod_{b \in \text{bin}} L_{ibr}^{bin} \prod_{b \in \text{ord}^o} L_{ibr}^{o}, \tag{G.11}$$

where bin indexes the marriage outcome and ord^o collects education, labor supply, and number of children.

Table G.1 summarizes the distributional assumptions for all first-stage dependent variables.

G.1.4 Overall first-stage likelihood

The overall first-stage likelihood, conditional on type r, is the product of the measurement-system and outcome-equation components:

$$L_{ir}^{(1)} = L_{ir}^{(m)} \cdot L_{ir}^{(o)}, \tag{G.12}$$

Table G.1: Summary of first-stage distributional assumptions

Variable	Distribution	Estimator	Categories
Measurements			
Age at first child	normal	normal MLE	—
Age at first relationship	normal, left-censored at $\underline{y} = 19$	normal MLE	—
Number of siblings	ordered categorical	ordered logit	0, 1, 2, 3+
Outcomes			
Education	ordered categorical	ordered logit	primary, secondary, bachelor, master/PhD
Marriage	binary	binary logit	partnered or not
Labor supply	ordered categorical	ordered logit	not working, part-time, full-time
Number of children	ordered categorical	ordered logit	0, 1, 2, 3+

Notes: The two timing measurements enter conditional on the event occurring within the observation window; a woman for whom it does not contributes a factor of unity for that measurement, and the extensive margin instead informs her type through the marriage and number-of-children equations. See the text for why this is not modeled as right-censoring.

where the superscript (1) indicates that $L_{ir}^{(1)}$ is the first-stage auxiliary likelihood.

G.1.5 EM algorithm

Given the first-stage likelihood in Equation (G.12), we estimate the posterior type probabilities q_{ir} and the market–cohort type shares $\pi_{r,k,c}$ via the EM algorithm. The E-step and M-step are iterated until convergence; we describe the t-th iteration of each.

E-step. At the t-th E-step, we treat the current parameter vector $\theta_{type}^{(t)}$ and type shares $\pi_{r,k,c}^{(t)}$ as given, and update the individual-level posterior as

$$\Pr(r|i) = q_{ir}^{(t)} = \frac{\pi_{r,k_i,c_i}^{(t)} L_{ir}^{(1)}\left(\theta_{type}^{(t)}\right)}{\sum_{r'} \pi_{r',k_i,c_i}^{(t)} L_{ir'}^{(1)}\left(\theta_{type}^{(t)}\right)}. \tag{G.13}$$

M-step. At the t-th M-step, we treat the posteriors $q_{ir}^{(t)}$ as known weights. Let $\mathcal{I}_{k,c}$ be the set of sampled women in market k and cohort c, with size $N_{k,c}$. The type shares are updated separately for each market–cohort cell:

$$\pi_{r,k,c}^{(t+1)} = \frac{1}{N_{k,c}} \sum_{i \in \mathcal{I}_{k,c}} q_{ir}^{(t)}, \tag{G.14}$$

and the parameters are updated by maximizing the weighted log-likelihood

$$\theta_{type}^{(t+1)} = \arg\max_{\theta_{type}} \sum_{i=1}^{N} \sum_{r=1}^{R} q_{ir}^{(t)} \log L_{ir}^{(1)}\left(\theta_{type}\right). \tag{G.15}$$

Because the type-specific intercept $\tilde{\gamma}_{br}$ is the only parameter that varies with r, and no parameter is shared across equations, the M-step separates into one weighted single-equation maximum-likelihood problem per dependent variable, so each M-step performs a full maximization rather than a partial update. We iterate until the observed-data log-likelihood $\ell^{(1)}$,

which increases monotonically across iterations, and the type shares $\pi_{r,k,c}$ have stabilized.

G.1.6 Second Stage Estimation

The first-stage estimates enter the second stage in two ways. For each market and cohort, $\pi_{r,k,c}$ determines the female fertility-type masses used to solve the marriage market equilibrium, as described in Appendix D. The posterior probabilities q_{ir} weight data moments indexed by fertility type. For example, let $S_i = 1$ if woman i is single and let E_i denote her education. The data moment for the share of single women with education e and fertility type r in market k and cohort c is

$$\widehat{m}^{\text{single}}_{e,r,k,c} = \frac{1}{N_{k,c}} \sum_{i \in \mathcal{I}_{k,c}} \mathbf{1}\{S_i = 1, E_i = e\}\, q_{ir}. \tag{G.16}$$

The denominator is the number of women in the market–cohort cell, as in the single-share moments described in Section 5.2.2. The same posterior weighting constructs matched-couple and child-number moments indexed by the wife's fertility type; the type shares enter the corresponding model moments through equilibrium.

H Revenue-Neutral Policy Exercises

The main counterfactuals do not rebate changes in tax revenue or charge households for policy costs. Here we re-solve each policy equilibrium with a per-adult lump-sum tax τ^* that restores baseline government net revenue, measured as income-tax revenue less policy outlays, tax expenditures, and existing transfers. Every adult pays the same tax, including both spouses in a couple. For the wage experiments, we treat the increase in women's earnings as a government-financed wage top-up. This assigns a cost to the wage changes within the model; it does not capture employer incidence or responses of labor demand, prices, and productivity. We assign the reduction in the deterministic education cost its consumption-equivalent monetary value and treat this amount as the public expenditure required per graduate to produce the cost reduction. Baseline income-tax revenue is 213 per unit of population, or 108 per adult, in the model's income units of DKK 1,000 per year. Table H.1 reports the revenue-neutral outcomes, and Figure H.1 plots the welfare effects. The results depend on this financing convention and the model's simplified tax-and-transfer system.

The balancing tax does not reverse any fertility effect, although it changes several magnitudes. Unconditional free childcare raises fertility by 0.060 children per woman, removing the motherhood wage penalty raises it by 0.028, and each education subsidy raises it by 0.023. The tax raises labor supply in several experiments through an income effect and also changes fertility through household income and marriage.

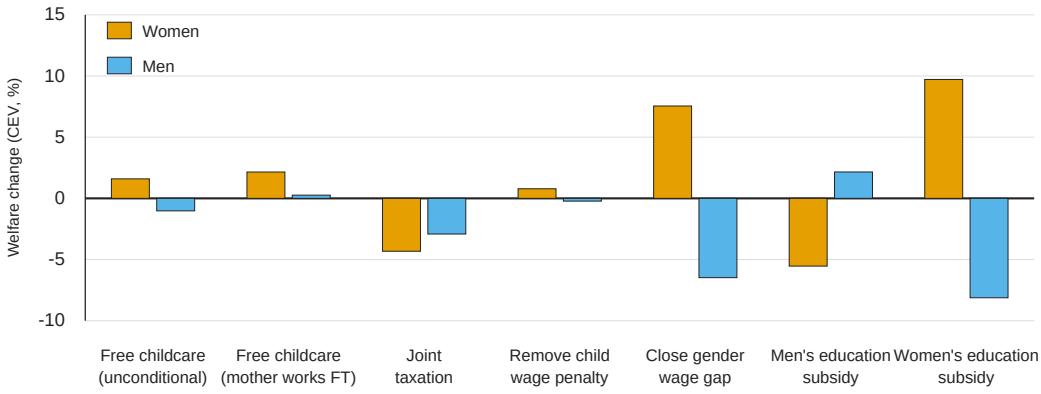
Behavioral responses recover 21 percent of the fiscal cost of unconditional childcare and 54 percent of the cost of work-contingent childcare. The work requirement reduces the direct cost from 19.0 to 16.4 per unit of population and encourages full-time work, increasing income-tax revenue. After financing, the full-time share among working women is 89.1 percent under work-contingent childcare, compared with 81.2 percent at baseline. After the balancing tax,

Table H.1: Revenue-neutral policy exercises at the later-cohort baseline

Policy	Fiscal cost	Lump-sum tax τ^*	Self-financing	Fertility	Welfare (CEV %) Women	Men
Family transfers						
Free childcare (unconditional)	19.0	7.6	21%	+0.060	+1.6	−1.0
Free childcare (mother works full-time)	16.4	3.8	54%	+0.043	+2.2	+0.3
Taxing the couple						
Joint taxation (income aggregation)	−12.8	0.4	n.a.	−0.036	−4.3	−2.9
Wage policies						
Remove motherhood wage penalty	13.8	2.7	61%	+0.028	+0.8	−0.2
Close gender wage gap	51.2	8.7	66%	−0.034	+7.5	−6.5
Education subsidies						
Men's education subsidy	49.9	13.8	45%	+0.023	−5.5	+2.2
Women's education subsidy	55.6	16.1	43%	+0.023	+9.7	−8.1

Notes: Each row re-solves the policy equilibrium under a per-adult lump-sum tax τ^* that returns income-tax revenue net of policy outlays, tax expenditures, and existing transfers to its baseline level. Every adult pays τ^*, so a couple pays twice that amount. "Fiscal cost" is the direct cost of the policy at its post-policy allocation, before accounting for induced changes in tax revenue and existing transfers; a negative value denotes a direct revenue gain. "Self-financing" is the share of a positive fiscal cost recovered through induced changes in income-tax revenue and existing transfers. It is not reported for policies with a negative fiscal cost. For the wage policies, fiscal cost treats the increase in women's earnings as a government-financed wage top-up. For the education policies, the reduction in deterministic education cost is assigned its consumption-equivalent monetary value, $c_e(e^{\Delta_e} - 1)$ at after-tax full-time income c_e, and this amount is treated as public expenditure per graduate. Fiscal cost is measured per unit of population (1.98 adults) and τ^* per adult, both in DKK 1,000 per year. "Fertility" is the change in completed fertility per woman, and welfare is the gender-specific consumption-equivalent change defined in (32); both include the effects of the financing tax.

Figure H.1: Welfare effects by gender under revenue-neutral policies



Notes: Gender-specific welfare changes in consumption-equivalent percent from the later-cohort baseline, with the full equilibrium re-solved under the per-adult lump-sum tax τ^* defined in the text. Figure 15 in Section 7.6 reports the corresponding effects without an offsetting fiscal adjustment.

unconditional childcare raises women's welfare by 1.6 percent and lowers men's by 1.0 percent; work-contingent childcare raises both, by 2.2 and 0.3 percent. Removing the motherhood wage penalty recovers 61 percent of its fiscal cost and raises women's welfare by 0.8 percent, with little effect on men's welfare. Joint taxation lowers both genders' welfare even though it mechanically raises revenue, because the resulting reductions in full-time work and partnership more than offset the initial revenue gain. Under the wage-top-up convention, closing the gender wage gap raises women's welfare by 7.5 percent and lowers men's by 6.5 percent.

The education subsidies imply public expenditures of 50 and 56 per unit of population,

Table H.2: Welfare effects by gender and education

	All		Women, by education				Men, by education			
Policy	Women	Men	Prim.	Sec.	Bach.	MA/PhD	Prim.	Sec.	Bach.	MA/PhD
Panel A: No offsetting fiscal adjustment										
Free childcare (unconditional)	+5.6	+3.0	+3.0	+5.5	+6.1	+5.7	+0.5	+3.1	+4.2	+3.7
Free childcare (mother works full-time)	+4.1	+2.2	+0.9	+3.8	+4.7	+5.1	+0.0	+2.2	+3.4	+3.3
Joint taxation (income aggregation)	−4.1	−2.7	+0.5	−2.7	−5.2	−9.1	+1.4	−1.4	−5.9	−9.1
Remove motherhood wage penalty	+2.2	+1.2	+0.1	+2.7	+3.0	+0.1	+0.4	+1.7	+1.7	−0.4
Close gender wage gap	+12.0	−1.9	+10.6	+11.5	+12.9	+12.3	−2.3	−2.0	−1.3	−1.9
Men's education subsidy	+1.6	+9.6	−6.8	−3.9	+6.3	+8.9	−2.8	−3.6	+35.1	+19.5
Women's education subsidy	+19.1	+0.2	−3.2	−2.4	+38.6	+24.2	−7.1	−5.4	+11.4	+9.7
Panel B: Revenue-neutral (lump-sum tax)										
Free childcare (unconditional)	+1.6	−1.0	−2.6	+1.0	+2.4	+3.5	−4.7	−1.4	+0.8	+1.8
Free childcare (mother works full-time)	+2.2	+0.3	−1.9	+1.6	+2.9	+4.0	−2.6	−0.0	+1.7	+2.4
Joint taxation (income aggregation)	−4.3	−2.9	+0.2	−2.9	−5.4	−9.2	+1.1	−1.7	−6.1	−9.2
Remove motherhood wage penalty	+0.8	−0.2	−2.0	+1.1	+1.6	−0.7	−1.5	+0.1	+0.5	−1.1
Close gender wage gap	+7.5	−6.5	+4.4	+6.4	+8.8	+9.9	−8.2	−7.1	−5.2	−4.2
Men's education subsidy	−5.5	+2.2	−16.5	−12.0	−0.9	+4.2	−12.2	−11.8	+26.4	+15.1
Women's education subsidy	+9.7	−8.1	−14.8	−12.0	+27.5	+17.9	−17.2	−14.6	+2.8	+4.4

Notes: Gender-specific consumption-equivalent welfare changes, in percent, from the later-cohort baseline. The "All" columns report consumption-equivalent changes calculated from the population-average inclusive values in (32). The by-education columns condition on entering the marriage market with a given education and integrate over partner choice, equilibrium Pareto weights, and household choices; the women's columns also average over fertility types. They include the deterministic cost of the stated education level but not its idiosyncratic cost shock. Under an education subsidy, the subsidized levels therefore include the direct cost reduction as well as the change in marriage market value. The "All" column is not a weighted average of the conditional columns because individuals can re-optimize education. Panel A reports the policy equilibria without an offsetting fiscal adjustment, and Panel B reports the revenue-neutral exercises in Table H.1.

compared with 19 for unconditional childcare. With financing, the men's subsidy raises men's welfare by 2.2 percent but lowers women's by 5.5; the women's subsidy raises women's welfare by 9.7 percent but lowers men's by 8.1. For the men's subsidy, the positive marriage market spillover to women is smaller than their share of the financing burden under this convention. Panel B of Table H.2 reports the corresponding welfare changes by education. The tertiary-educated of the subsidized gender retain the largest gains, whereas less-educated individuals of both genders bear most of the financing cost.