ECARES

The Cost of Children within Marriage: Intra-household Allocations of Time and Consumption

Laurens Cherchye
KU Leuven

Bram De Rock
ECARES, Université libre de Bruxelles

Khushboo Surana
University of York.

May 2026

ECARES working paper 2026-15

ECARES
ULB - CP 114/04
50, F.D. Roosevelt Ave., B-1050 Brussels BELGIUM
www.ecares.org

The Cost of Children within Marriage: Intra-household Allocations of Time and Consumption*

Laurens Cherchye† Bram De Rock‡ Khushboo Surana§

May 2026

Abstract

We use a structural revealed preference (RP) approach to study how fertility is associated with intra-household allocations and individual welfare within married couples. Building on a collective consumption framework, we jointly model material consumption and time use under heterogeneous preferences and Pareto-efficient household decisions, using marital stability to identify resource sharing from cross-sectional data. Applying the method to the 2023 wave of the US Panel Study of Income Dynamics (PSID), we find substantial gender inequality in intra-household allocations, with men consistently receiving more resources. These disparities widen sharply with fertility. Households with children exhibit lower material and time welfare for both spouses, but the burden falls disproportionately on mothers. Material resources shift moderately, whereas the time costs of children fall overwhelmingly on women: mothers experience large declines in leisure and increases in home production time, leading to sharply higher time-poverty risks and a clear worsening of the underlying welfare distribution. These patterns are especially pronounced among the highly educated and document a dimension of the cost of children that household-level measures miss.

Keywords: fertility, individual welfare, intra-household allocation, revealed preference, material and time poverty.

JEL codes: D12, D13, C14, J12.

*We thank participants of the CAPE seminar series at UCLouvain St-Louis Brussels for insightful discussion.

†Department of Economics, University of Leuven (KU Leuven). E-mail: laurens.cherchye@kuleuven.be. Laurens Cherchye gratefully acknowledges the Fund for Scientific Research - Flanders (FWO) and the Research Fund KU Leuven for financial support. He is also thankful to the Van Eeckhoudt-Ruttens family for supporting the 600 years of KU Leuven Chair on "Growing Up in Poverty".

‡ECARES, Université Libre de Bruxelles, Department of Economics, University of Leuven (KU Leuven) and Corvinus Institute for Advanced Studies, Corvinus University of Budapest. E-mail: bram.de.rock@ulb.be. Bram De Rock gratefully acknowledges the FWO and FNRS for their support.

§Department of Economics and Related Studies, University of York. Email: khushboo.surana@york.ac.uk.

1 Introduction

Fertility choices have major consequences for individual well-being within households. Children increase household needs and generate substantial demands for time-intensive home production, including childcare and domestic work. These demands reshape both material consumption and the allocation of time between market labor, leisure, and home production. Because these adjustments are often distributed unequally across spouses, individual welfare may decline even when total household resources remain unchanged. Standard poverty and welfare measures, which evaluate living standards at the *household* level, therefore risk masking substantial intra-household inequality. A growing body of evidence confirms that resource sharing is far from equal and that individual-level welfare assessments frequently diverge from household-level conclusions (e.g. Lise and Seitz, 2011; Dunbar, Lewbel, and Pendakur, 2013; Calvi, 2020).

These concerns are particularly relevant in the context of declining fertility. While macroeconomic and demographic forces are widely discussed (e.g. Goldin, 2021, 2025; Doepke, Hannusch, Kindermann, and Tertilt, 2023; Gobbi, Hannusch, and Rossi, 2026), the private costs of children borne by parents are central to fertility behavior. These costs are not purely monetary. They also involve time, specialization, and the distribution of adjustments within couples. If childbearing systematically reduces one partner's welfare, then these individual-level costs may shape fertility incentives even when aggregate household resources are sufficient. They also matter more broadly for the environment in which children are raised, since parental time and material allocations are core inputs into family life and child development.

A large literature in labor economics documents substantial "child penalties" in earnings, career progression, and labor force participation, showing that motherhood is associated with long-lasting reductions in women's labor market outcomes (e.g. Kleven, Landais, and Søgaard, 2019; Cortés and Pan, 2023; Kleven, Landais, and Leite-Mariante, 2025). Recent contributions, including Goldin (2024), provide comprehensive evidence on how income trajectories diverge sharply after childbirth. While this literature highlights the income costs of children, our focus is different and complementary. We study how fertility is associated with *individual welfare inside the household*, using measures based on time use and material consumption rather than earnings. Put differently, this paper does not model fertility choice directly. Instead, it measures how the material and time costs associated with children are distributed within marriage, and how that distribution maps into individual welfare.

This paper also fits into a broader structural literature that studies welfare at the level of individuals rather than households and emphasizes the joint allocation of time and money

within families (Cherchye, De Rock, and Vermeulen, 2026; Cherchye, De Rock, Vermeulen, and Gobbi, 2026). Our contribution is to bring that perspective to the cost of children within marriage. We build on the revealed preference framework of Cherchye, De Rock, and Surana (2026); however, in contrast to their focus, we concentrate on a sample of individuals in prime fertility ages and take fertility, rather than labor market participation, as the central organizing dimension. This allows us to ask a more specific question: how is parity associated with the division of material resources, leisure, and home-production burdens within marriage, and how do these reallocations translate into individual welfare and poverty risks? In this sense, the paper complements the child-penalty literature by moving from earnings effects to the intra-household distribution of the material and time costs of children.

More specifically, we study how fertility affects *intra-household allocations* and *individual welfare* using a structural collective model of household behavior in which spouses choose consumption and time use bundles in a Pareto-efficient manner (Chiappori, 1988, 1992). The distribution of household resources reflects relative bargaining positions, which depend on outside options in the marriage market. By integrating the collective model with Beckerian marriage-market restrictions (Becker, 1973, 1974), we impose marital stability to discipline feasible intra-household allocations. These stability conditions capture the opportunity costs that shape specialization within couples and connect directly to fertility, since the presence of children interacts with time allocation, earnings opportunities, and bargaining power.

Methodologically, we rely on the revealed preference (RP) framework of Cherchye, Demuynck, De Rock, and Vermeulen (2017), which delivers a fully nonparametric characterization of stable intra-household allocations from a single cross-section. Two features make this framework particularly suited to our question. First, it imposes no functional form assumptions on preferences, technologies, or sharing rules, allowing for heterogeneous preferences across individuals. Second, and crucially, Cherchye, De Rock, and Surana (2026) extend the framework to handle *non-working individuals* by identifying their unobserved wages as *shadow wages* within the stability system. This allows us to evaluate welfare differences associated with fertility even when labor-force participation adjusts.

We apply this framework to the 2023 wave of the U.S. Panel Study of Income Dynamics (PSID), focusing on adults aged 21–45 and constructing marriage markets at the state level with realistic age-based partner sets. We analyze heterogeneity across fertility (0, 1, or 2+ children) and education (low vs. high), and compute set-identified measures of intra-household resource allocation: Relative Individual Costs of Equivalent Bundles (RICEBs), Costs of Equivalent Bundles (CEBs), and the associated equivalent-consumption and leisure profiles. These allow us to examine how fertility is associated with intra-household sharing rules and how the resulting disparities appear in the distributions of *individual* material welfare

and leisure, with material and time poverty providing a useful summary of their lower-tail implications.

Our results highlight four main findings. *First*, intra-household inequality is substantial and systematically favors men across all fertility and education groups. Men receive higher RICEBs and larger intra-household resource shares, and they enjoy higher levels of equivalent consumption even before children are present. *Second*, these gender gaps widen with fertility. Moving from childless couples to couples with children, women experience declines in both material resources and, even more strongly, leisure, while home production burdens rise sharply. *Third*, these fertility-related reallocations are visible throughout the distributions of material welfare and leisure and are especially stark in the lower tail, where they translate into elevated material and time poverty risks for women. *Fourth*, education improves overall living standards for both men and women, but it does not eliminate the fertility gradient in welfare: even highly educated mothers experience sizable downward shifts in both consumption and leisure once children are present.

The remainder of the paper is organized as follows. Section 2 presents the structural framework. Section 3 describes the data and empirical implementation. Section 4 reports the empirical results on intra-household allocations and individual welfare. Section 5 concludes with policy implications.

2 Theoretical framework

Our empirical analysis builds on the revealed preference framework of Cherchye, Demuynck, De Rock, and Vermeulen (2017) and its extension to labor-market nonparticipation in Cherchye, De Rock, and Surana (2026). The framework links observed household behavior—material consumption, time allocation, wages, and prices—to the underlying bargaining positions of individuals in stable marriages. For our purposes, it serves two roles. First, it provides the restrictions that identify feasible intra-household allocations. Second, it delivers the money-metric welfare objects that we later take to the data.

In what follows, we first describe households, their consumption and time-use decisions, and the budget environment they face. We then introduce the marriage market and provide formal definitions of feasible matches and stable outcomes. Building on this structure, we characterize the revealed preference conditions for stability and explain how approximate stability and unobserved shadow wages enter the empirical formulation. Finally, we show how this framework yields set-identified measures of individual material welfare. Throughout, the model remains fully nonparametric and accommodates arbitrary heterogeneity in individual preferences.

We restrict attention to the welfare of the male and female partners within a couple and abstract from an explicit production technology for child welfare. When such a technology is modeled, it typically enters parental utilities in a weakly separable manner, as in Blundell, Chiappori, and Meghir (2005) and Cherchye, De Rock, and Vermeulen (2012). In our setting, monetary and time investments in children are instead absorbed into the household's public consumption and home-production components, consistent with the limited information on child-specific expenditures and time inputs in our data (see Section 3). This simplification preserves the RP content of the framework while allowing for heterogeneity in how parents value children's welfare.

2.1 Households, goods, and time

Consider a household composed of a man m and a woman w. The couple chooses a level of private material consumption $q_{m,w} \in \mathbb{R}_+$ and a level of public material consumption $Q_{m,w} \in \mathbb{R}_+$, where public consumption benefits both spouses jointly. Throughout, private and public material consumption are treated as Hicksian aggregates and their prices are normalized to unity. Private consumption is allocated between spouses, giving individual quantities $q^m_{m,w}$ and $q^w_{m,w}$ such that

$$q_{m,w} = q^m_{m,w} + q^w_{m,w}.$$

Time enters the household technology through leisure, market work, and home production. Each spouse $i \in \{m, w\}$ allocates a fixed time endowment T according to

$$T = \ell_i + o_i + h_i, \tag{1}$$

where ℓ_i denotes leisure, o_i denotes market labor, and h_i denotes home-production time. Home production generates a nonmarket public good. Following Cherchye, De Rock, and Surana (2026), we treat it as a one-input constant-returns technology, so that its value is summarized by the foregone wage associated with time spent in home production.

Wages (ω_m, ω_w) determine the opportunity cost of time. For individuals who do not participate in the labor market, wages enter as shadow wages inferred from the revealed preference conditions (see Section 2.4). We denote the couple's *full income* by $y_{m,w}$, defined as the sum of nonlabor income and the value of both spouses' time endowments evaluated at the relevant wages. In counterfactual states of singlehood, the corresponding full incomes $(y_{m,\emptyset}, y_{\emptyset,w})$ reflect the wages and nonlabor incomes associated with those scenarios and enter the stability constraints below. We return to the construction of these income variables in

Section 3.

2.2 Marriage market and stable matching

The marriage market consists of finite sets of men M and women W. A matching is a function

$$\sigma : M \cup W \to M \cup W$$

such that for all $m \in M$ and $w \in W$,

$$\sigma(m) = w \quad \Longleftrightarrow \quad \sigma(w) = m.$$

For notational simplicity, we do not explicitly consider singles. In our framework, single females (males) can be modeled as (virtual) couples with the male (female) consuming nothing.

A *matching allocation* specifies, for each couple $(m, \sigma(m))$,

$$\left(q^m_{m,\sigma(m)},\, q^w_{m,\sigma(m)},\, Q_{m,\sigma(m)},\, \ell_m, h_m,\, \ell_{\sigma(m)}, h_{\sigma(m)}\right).$$

A matching is *stable* if it satisfies two conditions:

1. **Individual Rationality (IR):** no spouse prefers being single at the wages and prices relevant for singlehood.

2. **No Blocking Pairs (NBP):** no unmatched pair (m, w) can jointly afford an allocation that makes both at least as well off as under σ, with at least one strictly better off, at the wages and prices relevant for their hypothetical match.

Although expressed in terms of utilities, the RP approach avoids imposing any functional-form assumptions. As we show next, both the IR and NBP conditions translate into linear inequalities involving only a limited number of unobserved variables. This characterization forms the basis of our empirical analysis. Importantly, as shown by Cherchye, Demuynck, De Rock, and Vermeulen (2017), the RP framework does not require any preference homogeneity assumptions and therefore accommodates fully heterogeneous individual preferences.

2.3 Revealed preference characterization

To describe the RP conditions under which the observed cross-sectional data can be interpreted as arising from a stable matching, we begin by summarizing the information contained in

the data. For now, we assume that all individuals are employed, so their wages are observed. In Section 2.4, we show how the RP conditions can be modified to account for unobserved wages among those not participating in the labor market. We assume that the data set D includes the following elements:

- the matching function σ;
- for every matched couple $(m, \sigma(m))$, the leisure levels $(\ell_m, \ell_{\sigma(m)})$, total private consumption $q_{m,\sigma(m)}$, and total public consumption $X_{m,\sigma(m)} = (Q_{m,\sigma(m)}, h_m, h_{\sigma(m)})$;
- for any matched or hypothetical couple (m, w), the full income $y_{m,w}$, the individual wages (ω_m, ω_w), and the public-good price vector $\Pi_{m,w} = (1, \omega_m, \omega_w)$;
- for any single man m, the income $y_{m,\emptyset}$ and individual wage ω_m; for any single woman w, the income $y_{\emptyset,w}$ and individual wage ω_w.

We observe aggregate public material consumption $Q_{m,\sigma(m)}$ and aggregate private material consumption $q_{m,\sigma(m)}$ only for matched couples $(m, \sigma(m))$. Since individual private consumption levels $q^m_{m,\sigma(m)}$ and $q^{\sigma(m)}_{m,\sigma(m)}$ are not observed, they must be inferred subject to the adding-up condition

$$q_{m,\sigma(m)} = q^m_{m,\sigma(m)} + q^{\sigma(m)}_{m,\sigma(m)}.$$

Next, for any potential marital outcome, we observe wages and incomes. For simplicity, we assume in this paper that labor productivity does not depend on marital status: individual wages ω_m and ω_w are taken to remain constant across all marital states, including remarriage or becoming single.[1] In Section 3, we return to the construction of the incomes $y_{m,w}$, $y_{m,\emptyset}$, and $y_{\emptyset,w}$ in our empirical application.

For any couple (m, w), we decompose the public-good price vector $\Pi_{m,w}$ into personalized prices $(\pi^m_{m,w}, \pi^w_{m,w})$, which represent each individual's willingness to pay for the publicly consumed quantities. These personalized prices can be interpreted as Lindahl prices, since they satisfy the adding-up restriction

$$\pi^m_{m,w} + \pi^w_{m,w} = \Pi_{m,w},$$

which reflects the requirement for Pareto-efficient provision of public goods. As with the individual private-consumption quantities, these personalized prices are unobserved and treated as unknowns in the RP conditions.

[1] This constant-productivity assumption follows Cherchye, Demuynck, De Rock, and Vermeulen (2017). As these authors discuss, it is in principle possible to relax this assumption within the revealed preference framework we adopt here. We abstract from doing so in order to keep the exposition simple.

These conditions are stated in the following result, which adapts Proposition 1 of Cherchye, Demuynck, De Rock, and Vermeulen (2017) to our setting.[2]

Proposition 1 (RP characterization of stable matching) *The data set D is rationalizable by a stable matching if and only if there exist individual private-consumption quantities $(q^m_{m,\sigma(m)}, q^w_{m,\sigma(m)})$ for all matched pairs $(m, \sigma(m))$ and personalized Lindahl prices $(\pi^m_{m,w}, \pi^w_{m,w})$ for all couples (m, w) satisfying the adding-up condition above, such that for every matched couple $(m, \sigma(m))$:*

(IR) ***Individual Rationality.*** *For every individual man m and woman w,*

$$y_{m,\emptyset} \;\leq\; q^m_{m,\sigma(m)} + \omega_m \ell_m + \Pi_{m,\sigma(m)} X_{m,\sigma(m)},$$

$$y_{\emptyset,w} \;\leq\; q^w_{\sigma(w),w} + \omega_w \ell_w + \Pi_{\sigma(w),w} X_{\sigma(w),w}.$$

(NBP) ***No Blocking Pairs.*** *For every unmatched pair (m, w),*

$$y_{m,w} \;\leq\; q^m_{m,\sigma(m)} + q^w_{\sigma(w),w} + \omega_m \ell_m + \omega_w \ell_w + \pi^m_{m,w} X_{m,\sigma(m)} + \pi^w_{m,w} X_{\sigma(w),w}.$$

These conditions have a clear economic interpretation. The IR inequalities require that each spouse's observed bundle—private consumption, leisure valued at the wage, and public consumption valued at the public good prices —be at least as affordable as what the spouse could obtain if single. The NBP inequalities ensure that no unmatched man and woman pair could jointly afford an allocation that makes both partners at least as well off as in their current match. Together, the IR and NBP conditions form a linear system in the unobserved individual private consumptions and personalized public-good prices. Feasibility of this system is equivalent to the observed behavior being rationalizable by a Pareto-efficient and stable matching. If the system is infeasible, then no stable marriage market consistent with the RP framework can generate the observed data. Operationally, rationalizability of a data set D reduces to checking the feasibility of a set of linear inequalities.

2.4 Approximate stability and shadow wages

In our empirical application, we use two refinements of the RP conditions in Proposition 1 when bringing them to the data. These refinements are discussed extensively in Cherchye, De Rock, and Surana (2026). We refer to that paper for formal details and provide only an intuitive overview here.

[2]Cherchye, Demuynck, De Rock, and Vermeulen (2017) establish the necessity of these RP conditions, while sufficiency is shown in Browning, Cherchye, Demuynck, De Rock, and Vermeulen (2024).

First, exact stability imposes sharp restrictions that real-world data need not satisfy literally. Small violations may arise from measurement error, frictions, non-economic determinants of partner choice, or idiosyncratic shocks. To address this, we introduce stability indices that enter the IR and NBP inequalities as additive relaxations. These indices measure the minimal deviation from exact rationalizability required to reconcile the data with stable matching. After adjusting for them, the resulting data set becomes exactly rationalizable.

Second, we refine the treatment of wages. While observed wages identify the opportunity cost of time for employed individuals, nonparticipants require shadow wages. Within the RP framework, these shadow wages are treated as additional unknown variables, restricted only by nonnegativity and the linear structure of the IR and NBP inequalities. They are therefore determined jointly with the allocations of private consumption and the personalized prices.

2.5 Individual welfare and set-identification

Having characterized the full set of intra-household allocations that are consistent with stable matching, we now turn to the evaluation of individual welfare. In our empirical analysis, individual welfare has two dimensions: material consumption and the allocation of time. Time use enters directly because leisure is treated as a privately consumed good and is observed in the data. This allows us to study time-related welfare outcomes, including time poverty, alongside material well-being.

To assess how households divide material resources between spouses, we rely on money-metric welfare measures that translate each spouse's within-marriage bundle into the expenditure that would be required if living alone. Following Cherchye, De Rock, and Surana (2026), we use three complementary concepts, each capturing a distinct aspect of individual welfare: the RICEB reflects each spouse's relative advantage, the intra-household shares describe the division of resources within the couple, and the CEB evaluates each spouse's absolute material standard of living. Taken together, these measures provide a compact picture of intra-household inequality.

Relative Individual Cost of Equivalent Bundle (RICEB). The RICEB measures a spouse's relative position within the household. It indicates the fraction of total household expenditure that an individual would need, as a single person facing market prices, to replicate the private and public goods consumed in marriage. For a couple $(m, \sigma(m))$, the RICEBs are defined as

$$R^m_{m,\sigma(m)} = \frac{q^m_{m,\sigma(m)} + Q_{m,\sigma(m)}}{q_{m,\sigma(m)} + Q_{m,\sigma(m)}}, \qquad R^{\sigma(m)}_{m,\sigma(m)} = \frac{q^{\sigma(m)}_{m,\sigma(m)} + Q_{m,\sigma(m)}}{q_{m,\sigma(m)} + Q_{m,\sigma(m)}}.$$

These values lie between zero and one. A higher RICEB means that an individual would require a larger share of total household expenditure to reproduce the bundle enjoyed in marriage, and therefore signals a more favorable relative position in the intra-household allocation.

Intra-household shares. While the RICEBs capture relative advantage, intra-household shares quantify the actual division of total household resources between spouses. They compare the cost of reproducing each spouse's within-marriage bundle, if single, to the counterfactual cost of reproducing both spouses' bundles as singles:

$$\gamma^{m}_{m,\sigma(m)} = \frac{q^{m}_{m,\sigma(m)} + Q_{m,\sigma(m)}}{q_{m,\sigma(m)} + 2Q_{m,\sigma(m)}}, \qquad \gamma^{\sigma(m)}_{m,\sigma(m)} = \frac{q^{\sigma(m)}_{m,\sigma(m)} + Q_{m,\sigma(m)}}{q_{m,\sigma(m)} + 2Q_{m,\sigma(m)}}.$$

These shares take values between zero and one and, unlike the individual RICEBs, add up to one by construction. Equal sharing yields one half for each spouse, whereas extreme inequality assigns all resources to one spouse and none to the other. Intra-household shares therefore provide a direct measure of how total resources are divided across partners, complementing the RICEB's focus on relative standing.

Cost of Equivalent Bundle (CEB). The CEB captures absolute material welfare. It measures the monetary amount each spouse would require, as a single individual, to purchase exactly the bundle enjoyed in marriage:

$$\rho^{m}_{m,\sigma(m)} = q^{m}_{m,\sigma(m)} + Q_{m,\sigma(m)}, \qquad \rho^{\sigma(m)}_{m,\sigma(m)} = q^{\sigma(m)}_{m,\sigma(m)} + Q_{m,\sigma(m)}.$$

Because the CEB is expressed in monetary units, it is comparable across individuals and groups and serves as the basis for our measure of material poverty.

A practical challenge is that these material welfare measures depend on individual private consumption, which is not directly observed. Instead, it must be inferred from the linear RP constraints in Proposition 1, with marital stability serving as the key identifying assumption. Consequently, each welfare measure is *set-identified* rather than point-identified.

For each spouse, we compute lower and upper bounds by solving two linear programs: one that minimizes the welfare measure subject to all RP constraints, and one that maximizes it. This yields all welfare levels consistent with stable intra-household behavior. In the next section, we describe how the PSID data are mapped into these structural objects so that the RP restrictions can be brought to the data.

3 Data and setup

The theoretical framework developed above yields a system of linear restrictions that we can bring to the data. Implementing these restrictions empirically requires careful construction of (i) the sample, (ii) the marriage markets within which individuals operate, and (iii) the material and time use aggregates that feed into the RP stability conditions.

We begin by introducing the 2023 PSID sample on which our empirical analysis is based. We then describe the construction of individuals' marriage markets, combining state-level and realistic age-based partner restrictions. Next, we explain how we aggregate material consumption into private and public Hicksian bundles and how we measure leisure, market work, and home production at the individual level. Our construction of individual-specific marriage markets and the material-consumption and time-use variables follows the procedures of Cherchye, De Rock, and Surana (2026). For compactness, we refer to that paper for the underlying motivation and focus here on the elements most directly relevant for the present application.

A measurement note is in order regarding how we treat the time and money that parents invest in their children. In the 2023 PSID, child-related expenditures are recorded at the family level, and there is no dedicated "money-on-children" measure. While detailed categories of childcare time use could, in principle, be observed in the Child Development Supplement (CDS), these data were not collected in the 2023 wave. Instead, we rely on aggregate measures of childcare time reported by parents in the main survey. Specifically, we treat parental time devoted to children as part of home production and incorporate child-related spending into the public material aggregate. This aligns the empirical implementation with the modeling choices in Section 2, where we abstract from an explicit production technology for children's welfare while allowing these inputs to operate through household public consumption and home production. Section 4.3 provides more specific details on the time-use variables and the measures of private and public material consumption used in our empirical application.

3.1 Sample and marriage market construction

Our empirical application uses microdata from the 2023 wave of the Panel Study of Income Dynamics (PSID). We restrict the sample to adult individuals aged 21–45, an age range in which labor supply, fertility, and partnership choices are particularly active, and in which most households consist of either single adults or couples without older dependents. After imposing these restrictions and excluding observations with missing demographic, consumption, or time-use information, the final sample comprises 3,127 households: 1,289 couples, 762 single men, and 1,076 single women.

Table 1 reports summary statistics for the main variables in our analysis. Several patterns stand out. Couples are, on average, slightly older than singles and also display higher wages. Time allocation differs markedly across household forms: singles devote substantially more time to leisure, whereas within couples there is a pronounced gender gap in housework time. Both public and private consumption are considerably higher for couples, reflecting household size and the presence of shared goods. Finally, education and employment rates are high across all groups, though singles—particularly single men—have somewhat lower education shares and are much more likely to have no children. Together, these statistics illustrate the demographic and economic heterogeneity that our structural analysis seeks to capture.

Table 1: Summary statistics

	Couples		Singles	
	Male	Female	Male	Female
Age	36.13	34.86	33.30	34.31
Wage	37.36	32.57	27.96	25.28
Leisure	53.99	50.52	64.40	61.27
Housework	17.62	29.67	5.09	10.44
Public consumption	498.00	498.00	274.51	290.42
Private consumption	869.77	869.77	508.26	536.30
Fraction with high education	66.02	74.40	56.96	65.61
Fraction employed	94.96	84.10	87.53	86.71
Fraction with no children	29.87	29.87	85.43	50.74
Fraction with one child	22.50	22.50	8.40	20.91
Observation	1,289	1,289	762	1,076

Notes: Wage in dollars per hour, leisure and housework in hours per week, and consumption in dollars per week.

Marriage markets are defined at the *state* level, reflecting the idea that individuals are most likely to search for partners within their geographic region. Some U.S. states contribute relatively few observations to the PSID; to avoid placing individuals in unrealistically thin markets, we merge states with fewer than 50 observations into adjacent or demographically similar clusters. This yields 26 state-level marriage markets. Appendix A reports the sample size by state. Figure 1 presents a map of the resulting 26 marriage markets.

Within each geographic market, we construct individual-specific marriage markets using empirically grounded *age-difference* rules. Specifically, for each market, we allow men to consider potential female partners whose ages fall within bounds defined by the 2.5th and 97.5th percentiles of the observed age differences between men and women in couples in that market. These market-specific age brackets capture 95% of observed couples and ensure that feasible matches reflect realistic demographic opportunities.

Figure 1: Marriage markets

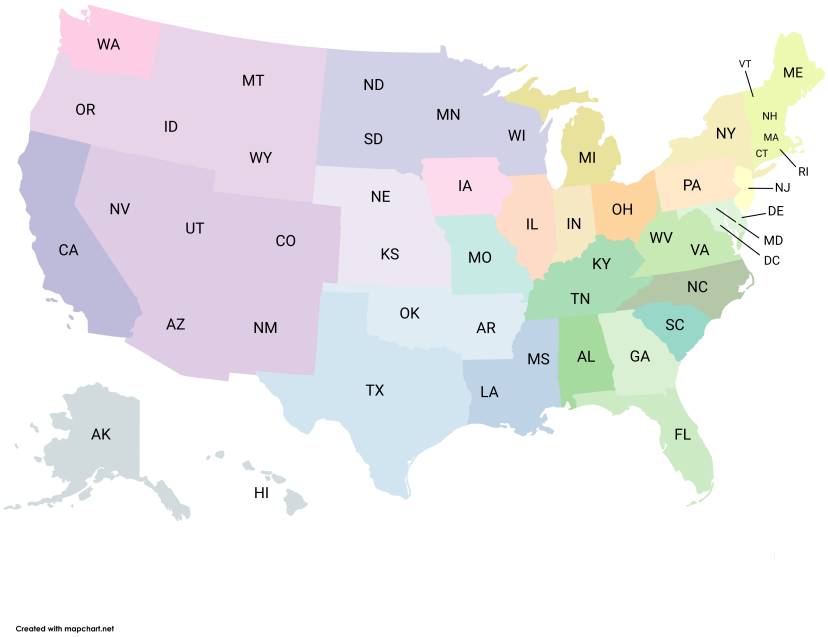

3.2 Material consumption, time use and household income

To empirically apply the collective model described in Section 2.1, we need measures of private and public material consumption, leisure, home production, and full household income.

Material consumption. We use Hicksian aggregates of private and public consumption, with prices normalized to unity. The private component comprises:

- food and drinks (at home and outside the home),
- schooling-related expenditures,
- computers, clothing, and recreation,
- 50% of expenditures on vacation, housing, transportation, childcare, and healthcare.

The remaining 50% of these latter categories is assigned to *public* material consumption. This classification reflects the fact that many expenditure categories contain both an individually consumed component and a shared component that benefits both spouses.

Time use. Time use is measured in terms of market work, leisure, and home production. Leisure is treated as a *private* good, consistent with the idea that it directly enters individual

utility. To measure leisure, we assume a daily time endowment of 16 hours, with the remaining 8 hours reserved for sleep and meals. Leisure is then computed as this time endowment minus reported hours of market work and housework. Home production is modeled as a *public* good because it generates domestic services that benefit both spouses. This category includes childcare and other child-related tasks, which constitute a substantial share of domestic work in households with children. As outlined in Section 2.1, time spent on leisure and home production is valued at its opportunity cost, as captured by the individual's wage.

Household income. Full household income consists of potential labor income and non-labor income. Potential labor income is defined as the value of the household's total time endowment evaluated at market wages, corresponding to a hypothetical scenario in which all available time (i.e. 16 hours per day) is devoted to market work.

For observed couples, non-labor income is derived using a consumption-based approach: it is calculated as total household expenditures minus labor income, where labor income equals the product of observed market hours and individual wages. For unobserved marital outcomes, non-labor income is defined as the sum of the spouses' individual non-labor incomes in their observed match.

Because non-labor income is constructed only at the household level for observed couples, individual non-labor incomes are not separately observed. We therefore treat them as unknown components that must sum to the household-level non-labor income implied by expenditures. To avoid implausibly extreme imputations, we impose a regularity condition requiring each spouse's non-labor income to lie between 40% and 60% of total household non-labor income. This restriction rules out corner solutions in which one spouse receives nearly all non-labor income, as well as imposing the equally unrealistic case of an exactly equal split following marital dissolution.

Summary statistics. As noted in the Introduction, our analysis focuses on the individual welfare effects associated with fertility while also allowing for differences by education. To this end, we classify individuals along two dimensions:

- **education**: low (high school or less) vs. high (more than 12 years),
- **number of children**: 0, 1, or 2+.

This classification applies to both singles and couples and serves as the main organizing structure for our empirical results.

Tables 2 and 3 report summary statistics by (education × children) type for men and women separately. The statistics include sample sizes, employment rates, leisure hours,

labor hours, material consumption, and average age. Several patterns emerge. Employment rates are uniformly high among men (above 80% for all low-educated types and above 95% for high-educated types), while women exhibit a modest decline in employment as fertility increases. For women, both leisure and labor hours fall in the presence of children. For men, leisure declines while labor hours increase, particularly among low-educated types. Consumption levels increase both with education and with the number of children, reflecting higher household needs as well as the greater earnings capacity of more educated individuals. A central question for our empirical analysis is how fertility affects the way this aggregate household consumption is allocated between spouses.

Table 2: Summary statistics by type: low education

kids	N	% empl.	leisure	labor	consumption	age
Panel A: males (low-educated)						
0	387	80.36	69.49	32.25	738.97	33.16
1	130	90.00	55.32	38.14	1065.13	34.95
2+	249	89.16	53.38	37.64	1157.49	36.88
Panel B: females (low-educated)						
0	220	77.73	71.96	29.45	784.43	32.24
1	153	76.47	53.91	28.27	876.63	34.29
2+	327	71.25	55.57	25.58	948.09	34.63

Table 3: Summary statistics by type: high education

kids	N	% empl.	leisure	labor	consumption	age
Panel A: males (high-educated)						
0	649	96.92	63.35	40.90	1051.67	33.09
1	224	95.54	49.29	38.85	1408.88	36.87
2+	412	96.60	46.46	41.06	1574.59	38.01
Panel B: females (high-educated)						
0	711	94.66	65.94	37.62	1037.41	32.54
1	362	88.40	46.31	33.32	1263.67	36.43
2+	592	84.97	42.47	31.13	1420.33	36.93

4 Empirical welfare analysis

We apply the RP identification strategy from Section 2 to the 2023 PSID sample introduced in Section 3 in order to study how fertility and education are associated with individual well-being through material consumption and time use. We first examine how the structural welfare objects introduced in Section 2.5—RICEBs, intra-household resource shares, equivalized

consumption, and leisure—vary across fertility and education groups. We then complement these descriptive patterns with regression evidence to see how this is associated with observed demographics. Finally, we present cumulative distribution functions (CDFs) of equivalized consumption and leisure together with the corresponding relative poverty thresholds. This allows us to assess not only whether fertility is associated with lower welfare on average, but also how it affects the lower tail of the welfare distribution. In that setting, material and time poverty provide useful summary measures, while the CDFs reveal whether the underlying shifts are local or more pervasive.

4.1 Descriptive patterns of intra-household inequality

Table 4 reports RP-identified bounds for the RICEB by gender, fertility (0, 1, or 2+ children), and education (low vs. high). Recall that a higher RICEB indicates that an individual would require a larger share of total household expenditure to replicate the within-marriage bundle if single, and therefore reflects a more favorable intra-household allocation.

The most robust pattern is the gender gap itself. Across all fertility–education groups, men have higher RICEBs than women, indicating that they systematically command a larger share of household resources. Variation across fertility groups is more modest and less monotonic than the gender difference, especially at the transition from zero to one child. Still, at higher parities the gap tends to widen, particularly among the highly educated, which is consistent with a shift in material advantage toward men in households with more children.

Table 4: RICEBs by fertility and education

Children	**Male**				**Female**			
	Low Edu.		High Edu.		Low Edu.		High Edu.	
	Lower	Upper	Lower	Upper	Lower	Upper	Lower	Upper
0	0.67	0.77	0.68	0.81	0.53	0.62	0.56	0.69
1	0.68	0.78	0.65	0.78	0.54	0.66	0.59	0.71
2+	0.71	0.79	0.73	0.83	0.52	0.60	0.55	0.65

Notes: Bounds reflect set-identification under RP stability conditions.

RICEBs capture individuals' relative advantage within the household in terms of material good consumption, but they do not directly quantify each spouse's share of total household resources. Table 5 therefore reports bounds on intra-household resource shares. These patterns closely mirror the RICEB results: men consistently receive larger shares, and the fertility-related changes in the distribution of material resources are modest relative to the large and systematic gender differences.

Table 5: Intra-household resource shares by fertility and education

Children	**Male share**				**Female share**			
	Low Edu.		High Edu.		Low Edu.		High Edu.	
	Lower	Upper	Lower	Upper	Lower	Upper	Lower	Upper
0	0.50	0.57	0.50	0.60	0.39	0.46	0.41	0.51
1	0.51	0.58	0.48	0.57	0.40	0.49	0.43	0.52
2+	0.53	0.59	0.54	0.61	0.39	0.45	0.40	0.47

Notes: Bounds reflect set-identification under RP stability conditions.

To complement these share-based measures of material consumption, Table 6 reports bounds on each individual's equivalized consumption together with observed leisure. Equivalized consumption is constructed from the CEB defined above. Because individuals differ in the number of children in their household, we adjust consumption for household size by assigning a weight of 0.5 to each child, following the original OECD equivalence scale. Operationally, this amounts to dividing the sum of individual private consumption and household public consumption by $1 + 0.5\times$ (number of children).

Three broad patterns emerge. First, across all fertility–education groups, men have higher equivalized consumption than women. Second, while higher fertility is associated with lower equivalized consumption for both spouses, the decline is larger and more systematic for women, especially at higher parities. Among the low-educated, women experience a steeper decline than men when crossing the zero-to-one child threshold, whereas for the high-educated, there is a larger consumption drop for men. However, among both education types, with two or more children, women experience a disproportionately larger reduction than their male counterparts.

Third, the leisure results reinforce this pattern: higher fertility is associated with lower leisure for both genders, but women lose substantially more leisure than men, with the largest declines among the highly educated. When moving from one to two or more children, high-educated men and women both experience a further decrease in leisure, reflecting a stronger preference for time investment in children among high-educated parents (Guryan, Hurst, and Kearney, 2008; Cortés and Pan, 2023).

Taken together, the descriptive evidence reveals substantial and systematic intra-household inequality. Material consumption is consistently distributed in favor of men, although the fertility-related differences in material allocations are relatively modest. In contrast, fertility is associated with a strong impact on the allocation of time, redistributing it in ways that disproportionately disadvantage women, with the largest effects observed among those with higher education. These patterns motivate the regression analysis that follows and highlight

Table 6: Equivalent consumption and leisure by fertility and education

Children	Equivalent consumption				Leisure (hrs/week)	
	Low Edu.		High Edu.		Low Edu.	High Edu.
	Lower	Upper	Lower	Upper	Level	Level
Males						
0	701.59	800.69	890.07	1053.06	69.44	62.31
1	525.35	592.31	658.17	776.82	54.84	49.06
2+	362.71	404.03	541.76	616.50	54.03	45.84
Females						
0	540.21	624.73	725.88	887.81	68.90	65.48
1	373.35	441.80	588.82	700.98	51.26	42.20
2+	257.84	292.52	399.47	472.35	53.46	39.61

Notes: Bounds reflect set-identification under RP stability conditions. Leisure is observed directly in the data and is therefore not bounded.

the importance of assessing welfare at the individual rather than the household level.

4.2 Regression evidence on the sources of inequality

To examine these patterns more systematically, Tables 7 and 8 report regressions that relate the four welfare indicators (RICEBs, intra-household resource shares, equivalized consumption, and leisure) to fertility, education, demographic characteristics, and employment for married men and women. Clearly, these regressions should be read as conditional associations rather than causal estimates of the effect of childbirth.

For married men (Table 7), the results broadly reinforce the descriptive patterns. Higher parity is not significantly associated with men's RICEBs or intra-household resource shares, suggesting that their relative position within the household changes little across fertility groups once other observables are held fixed. By contrast, the associations with absolute outcomes are large and precisely estimated: equivalized consumption is substantially lower for men, and leisure declines by about 21 hours per week for fathers with one child and by a similar amount for those with two or more children. Employment strongly raises men's equivalized consumption and resource shares, while mechanically reducing leisure. The interaction terms suggest that higher education only partly attenuates the fertility-related declines in men's absolute well-being.

For women (Table 8), the patterns are markedly different. Holding other observables fixed, women with one child have significantly higher RICEBs and intra-household shares than otherwise similar childless women, indicating a modest improvement in their relative position within the household. However, this relative gain does not translate into higher absolute

Table 7: Regression results: married males

	RICEB	**Intra-household share**	**Equivalent consumption**	**Leisure**
High education	0.0114	0.00604	164.1***	-4.162**
	(0.0244)	(0.0183)	(44.94)	(1.701)
One child	0.000621	-0.00347	-227.1***	-20.80***
	(0.0332)	(0.0249)	(50.91)	(2.830)
Two+ children	0.0256	0.0174	-422.5***	-20.57***
	(0.0268)	(0.0201)	(39.90)	(2.253)
High education × One child	-0.0348	-0.0285	-85.25	1.532
	(0.0369)	(0.0276)	(61.44)	(3.066)
High education × Two+ children	0.00369	-0.00145	-34.00	-0.715
	(0.0304)	(0.0228)	(49.72)	(2.447)
Age of youngest child	0.000314	0.000983	1.058	0.914***
	(0.00178)	(0.00132)	(2.800)	(0.159)
Employed	0.275***	0.207***	310.2***	-33.41***
	(0.0249)	(0.0179)	(27.71)	(2.807)
Age	0.000840	0.000576	4.824**	0.0449
	(0.00124)	(0.000916)	(2.054)	(0.101)
Home owner	0.00131	-0.00192	173.5***	-2.865**
	(0.0142)	(0.0106)	(21.48)	(1.190)
Observations	1289	1289	1289	1289

Notes: The first two columns (RICEB and intra-household share) are estimated via interval regression, reflecting the set-identified nature of the dependent variables. The last two columns are estimated by OLS; leisure is directly observed, while equivalized consumption is based on a point summary of the identified interval. Robust standard errors are reported in parentheses. Low education is defined as ≤ 12 years of schooling, while high education is defined as > 12 years. The reference group consists of individuals with low education, no children living in the household, and who do not own their home. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

welfare. Equivalized consumption is lower by more than 130 dollars per week for mothers of one child and by almost 280 dollars per week for mothers with two or more children, while leisure falls by roughly 26 to 30 hours per week. Higher education raises women's equivalized consumption but leaves their relative position within the household largely unchanged, and it is associated with an even steeper fertility-related reduction in leisure at higher parities.

Table 8: Regression results: married females

	RICEB	**Intra-household share**	**Equivalent consumption**	**Leisure**
High education	0.0179	0.0126	151.3***	-0.186
	(0.0256)	(0.0186)	(43.70)	(2.162)
One child	0.0787**	0.0552**	-131.1***	-30.33***
	(0.0373)	(0.0270)	(49.56)	(3.612)
Two+ children	0.0394	0.0281	-279.0***	-26.21***
	(0.0287)	(0.0208)	(40.96)	(2.929)
High education × One child	0.00465	-0.0000321	-4.984	-3.303
	(0.0400)	(0.0289)	(58.50)	(3.661)
High education × Two+ children	0.000176	-0.00492	-69.41	-9.249***
	(0.0321)	(0.0233)	(47.39)	(3.044)
Age of youngest child	-0.00474***	-0.00274**	-6.042***	1.366***
	(0.00174)	(0.00126)	(2.294)	(0.181)
Employed	0.215***	0.159***	220.0***	-20.24***
	(0.0145)	(0.0103)	(17.30)	(2.144)
Age	-0.0000534	-0.0000846	5.082***	0.0988
	(0.00125)	(0.000914)	(1.840)	(0.110)
Home owner	-0.0469***	-0.0364***	83.29***	-1.969
	(0.0144)	(0.0105)	(20.03)	(1.366)
Observations	1289	1289	1289	1289

Notes: The first two columns (RICEB and intra-household share) are estimated via interval regression, reflecting the set-identified nature of the dependent variables. The last two columns are estimated by OLS; leisure is directly observed, while equivalized consumption is based on a point summary of the identified interval. Robust standard errors are reported in parentheses. Low education is defined as ≤ 12 years of schooling, while high education is defined as > 12 years. The reference group consists of individuals with low education, no children living in the household, and who do not own their home. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Overall, the regression evidence reinforces the gender asymmetries identified in the descriptive analysis. Higher parity is associated with lower material consumption and less leisure for both men and women, but the magnitude of these associations is substantially larger for women. Moreover, men's relative positions within the household remain largely

unchanged across fertility groups, whereas women's RICEBs and resource shares rise only modestly and do not offset the large declines in their absolute living standards.

Appendix B shows that singles also exhibit lower consumption and leisure in the presence of children, again particularly among women. Unlike married individuals, however, singles do not have access to intra-household mechanisms through which some of the costs associated with children can be shared. Comparing singles and married individuals therefore suggests that some fertility-related inequalities, especially those involving the division of time, may be amplified within marriage, whereas others, such as the education gradient in consumption, operate more broadly.

4.3 Material and time poverty

We next examine individual-level poverty along two dimensions: material poverty, based on equivalized consumption, and time poverty, based on leisure. Both concepts follow a relative definition. Individuals are classified as materially poor if their equivalized consumption falls below 60% of the median, and as time poor if their leisure falls below 60% of the sample median.[3]

At the same time, the poverty thresholds should be seen as summary measures of broader distributional patterns. To assess whether the fertility-related differences documented in the previous subsection are confined to the lower tail or instead reflect more pervasive changes in individual welfare, we complement the poverty analysis with cumulative distribution functions (CDFs) of equivalized consumption and leisure by fertility and education group. The CDFs therefore remain central, while the poverty thresholds sharpen the lower-tail implications of the distributions.

Material poverty and the distribution of equivalized consumption. Figures 2 and 3 plot CDFs of equivalized consumption for married men and women across education–fertility groups. Material poverty rises with fertility for both genders, but the increase is much stronger for women. Among women, the CDFs shift upward as fertility rises, especially for those with two or more children, indicating that the deterioration is not confined to the bottom tail of the distribution. This implies not only higher material-poverty risks, but also a broader worsening of women's material welfare. For men, equivalized consumption also tends to decline with fertility, but the shifts are smaller and less uniform.

[3]Although this is a simple measure of time poverty, it effectively serves our purpose of illustrating the time-allocation effects associated with fertility. In practice, time poverty is a multidimensional concept that reflects insufficient time to meet personal and social needs and ideally requires attention to both the quantity and quality of time across domains of life; see Williams, Masuda, and Tallis (2016) for a more detailed discussion.

Education clearly improves material living standards, but it does not eliminate the fertility-related disparities. Within education groups, the gender differences largely mirror those documented in Tables 6, 7, and 8. In particular, even highly educated mothers experience substantial downward shifts in the distribution of equivalized consumption relative to childless women. The poverty threshold therefore helps quantify the lower-tail consequences of a broader result: fertility is associated with an adverse shift in women's material welfare, and this shift is strong enough to generate markedly higher material-poverty risks.

Time poverty and the distribution of leisure. Figures 4 and 5 report the corresponding CDFs for leisure. Here the patterns are even sharper. Time poverty increases strongly with fertility for both genders, but much more so for women. For mothers, and especially for those with two or more children, the leisure CDFs shift downward markedly, implying substantially higher time-poverty risks. In this case, the poverty threshold is particularly informative because it quantifies a pattern that is also visible in the full distribution: higher fertility is associated with a broad and substantial decline in women's leisure. Among men, the decline in leisure is also clearly visible, but it is much smaller, and the distributions rarely display the same degree of dominance observed for women.

Within education groups, the gender differences in time poverty remain large. If anything, the fertility-related decline in leisure appears particularly strong among the highly educated, suggesting that the time costs associated with children remain substantial even in households with higher material resources. These figures therefore indicate not only that mothers are more likely to fall below the time-poverty threshold, but also that the entire leisure distribution for mothers shifts in an adverse direction as fertility rises. This reinforces the message of the previous subsection: while fertility is associated with lower material welfare for women, its most pronounced welfare consequences operate through time.

Taken together, Figures 2–5 show that higher fertility is associated with downward shifts in both equivalized consumption and leisure, but that these shifts are much stronger for women than for men. The poverty thresholds make this especially transparent at the lower tail of the distribution. In that sense, poverty is not a separate idea from the distributional analysis, but a useful way of quantifying its sharpest welfare implications. The broader message remains that fertility is associated with a systematic worsening of individual welfare distributions, especially through the time burden borne by mothers.

Figure 2: CDF of log consumption by number of children and education; married males

Low education

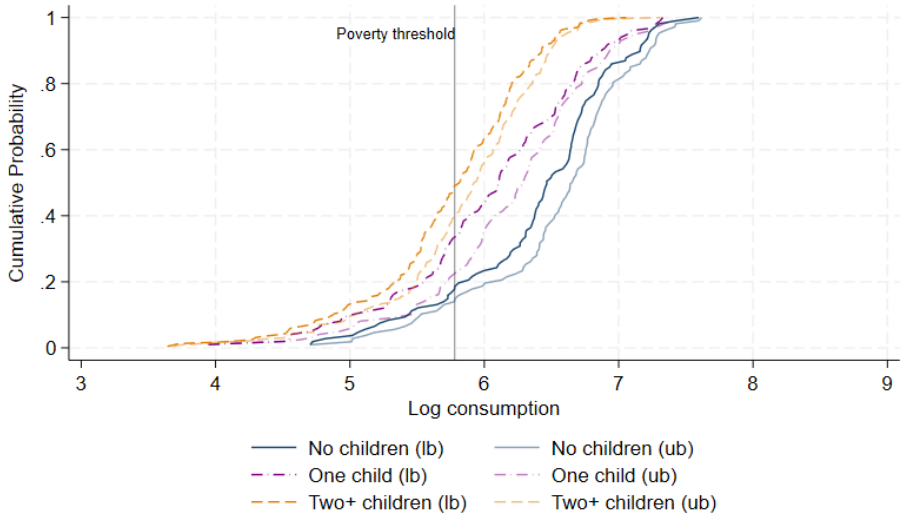

High education

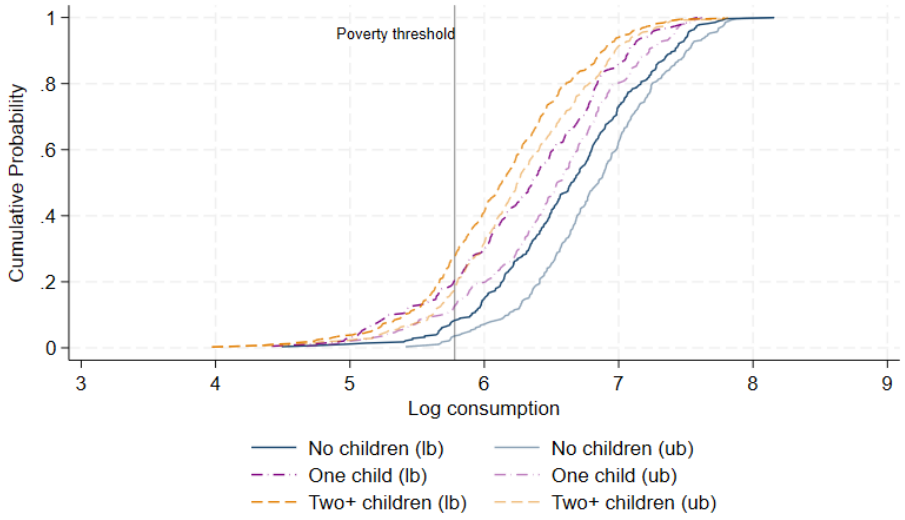

Notes: The consumption poverty threshold is set at 60% of the median equivalized consumption in the sample. For married individuals, whose consumption is set-identified, the figures are based on the midpoint of each individual's identified consumption interval.

Figure 3: CDF of log consumption by number of children and education; married females

Low education

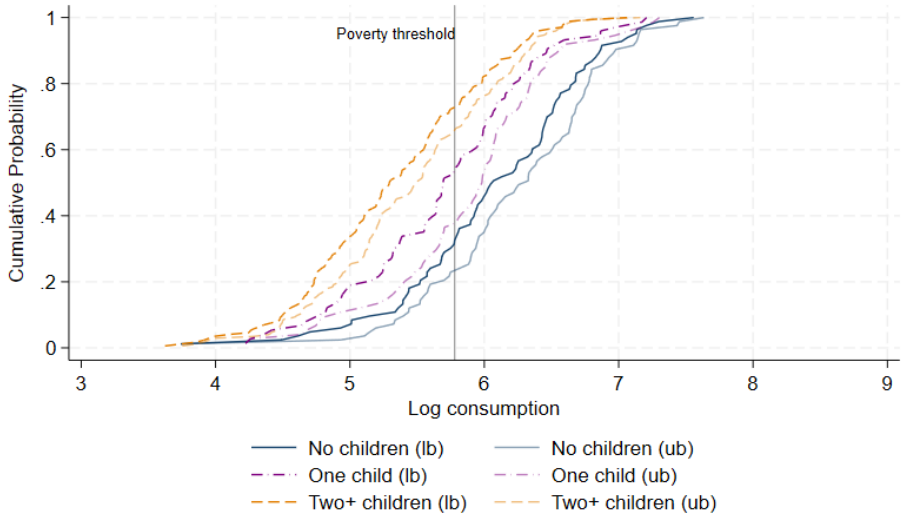

High education

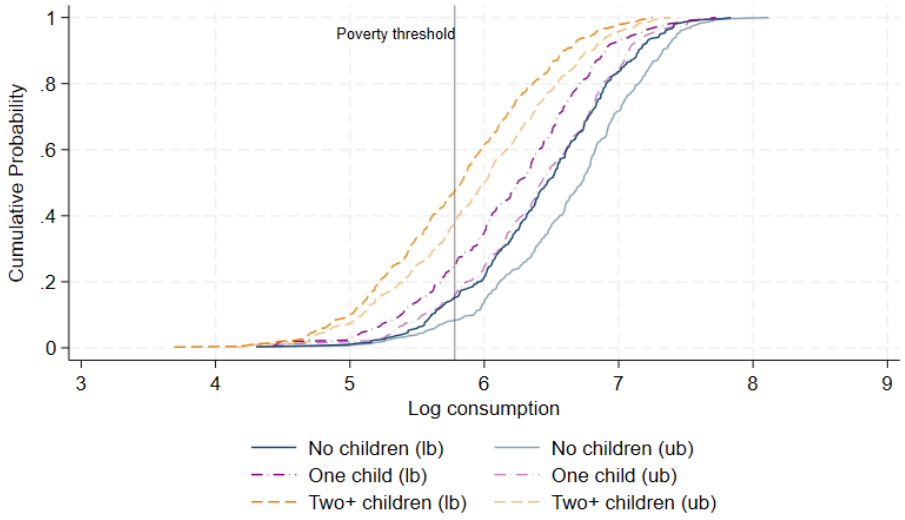

Notes: The consumption poverty threshold is set at 60% of the median equivalized consumption in the sample. For married individuals, whose consumption is set-identified, the figures are based on the midpoint of each individual's identified consumption interval.

Figure 4: CDF of leisure by number of children and education; married males

Low education

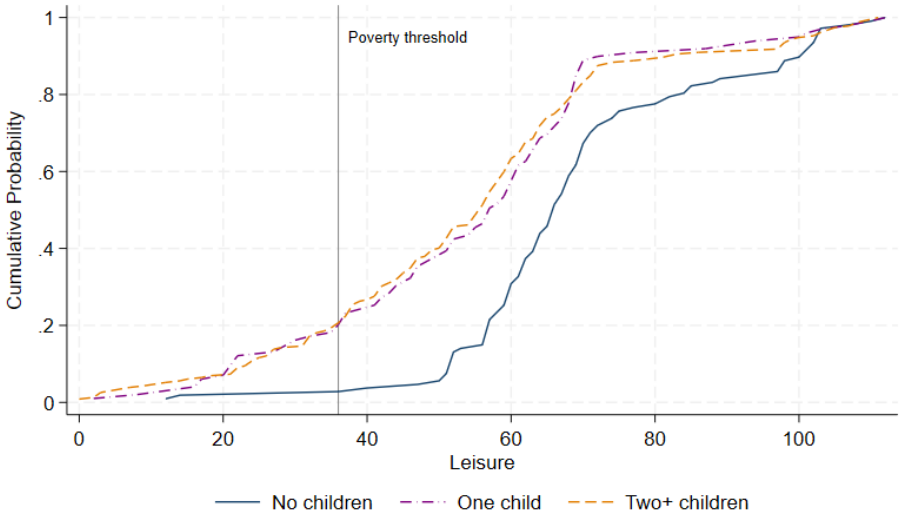

High education

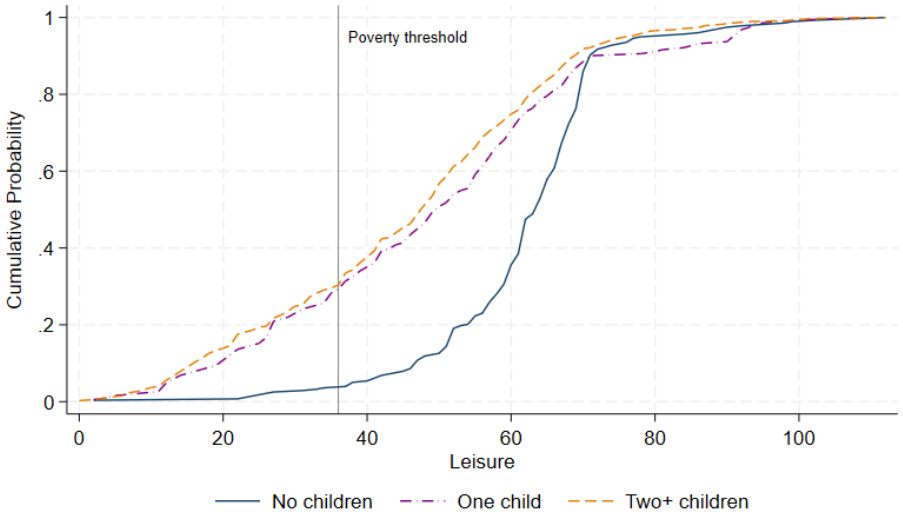

Notes: The time-poverty threshold is defined as 60% of the median individual leisure in the sample.

Figure 5: CDF of leisure by number of children and education; married females

Low education

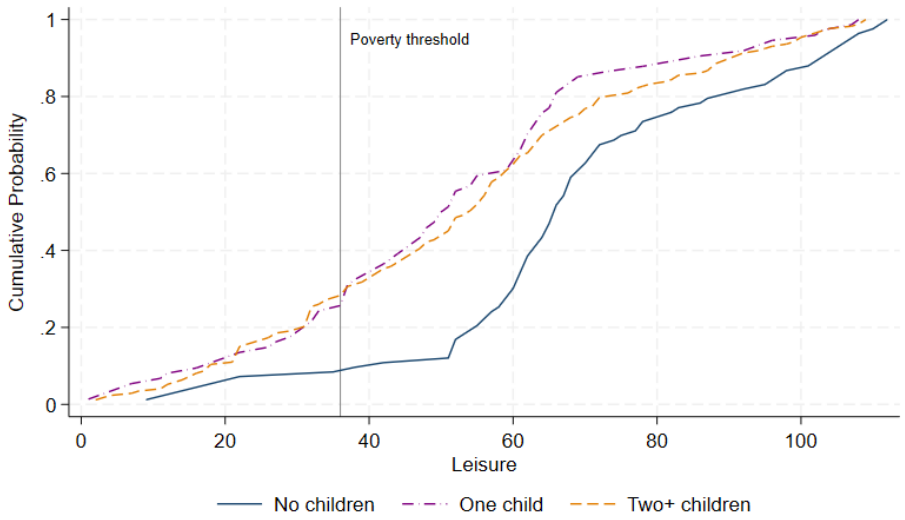

High education

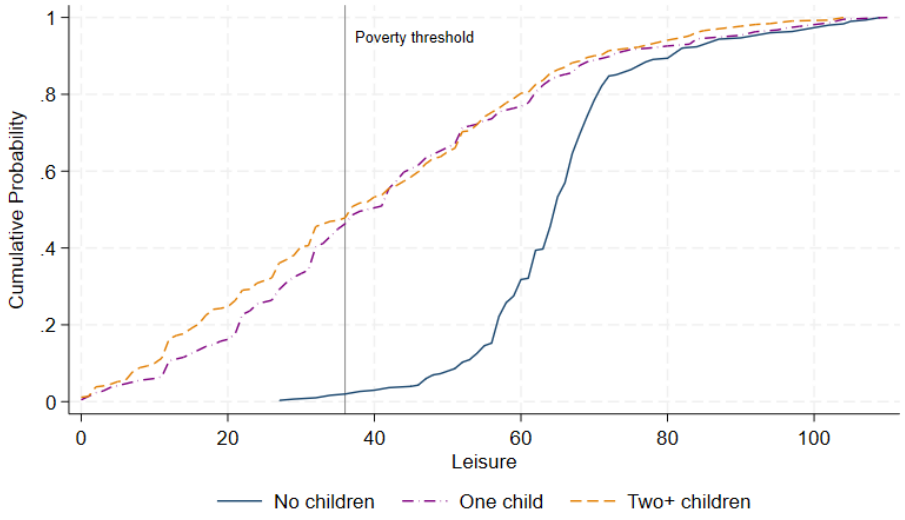

Notes: The time-poverty threshold is defined as 60% of the median individual leisure in the sample.

5 Discussion and conclusion

Our empirical analysis shows that fertility is associated with a pronounced reallocation of both money and time within marriage. Households with children, and especially those with two or more children, display larger gender gaps in equivalent consumption, leisure, and home-production burdens than childless couples. Men also experience declines in well-being as fertility rises, but these declines are considerably smaller than those observed for women. The central message is therefore not simply that children are costly, but that these costs are distributed unevenly within households. This creates a wedge between household-level resources and individual-level welfare that standard household measures do not capture.

These disparities are already visible among childless couples, but they become substantially more pronounced across fertility groups. The material dimension matters, yet the time dimension is even more consequential. Mothers experience large reductions in leisure and substantial increases in home production, whereas fathers' time allocations adjust much more modestly. This is why the time-poverty results are sharper and more pervasive than the material-poverty results. From this perspective, the paper complements the child-penalty literature by moving from earnings and labor-market outcomes to the intra-household distribution of the material and time costs associated with children. More broadly, it operationalizes a structural agenda that studies fertility through household behavior, the value of partnership, and the joint allocation of time and money within families (Cherchye, De Rock, Vermeulen, and Gobbi, 2026; Cherchye, De Rock, and Vermeulen, 2026).

Material poverty follows a similar but less pronounced pattern. Women, especially those with lower education, face substantially higher material poverty risks once children are present. Higher education improves overall living standards, but it does not eliminate the fertility-related disparities: even highly educated mothers experience sizable downward shifts in both consumption and leisure. Taken together, these findings suggest that fertility is associated not only with higher household needs, but also with a more unequal internal distribution of the resulting burdens.

These findings naturally point to policies that address both the material and time costs of children. Expanding access to affordable, high-quality childcare can directly reduce the time burden on mothers and shift part of home production to the market sector. Measures that promote more equal parental leave, flexible work arrangements, or greater paternal involvement can further counteract the specialization that often emerges after childbirth. Income-support policies, such as child benefits or subsidies for child-related goods, can ease material constraints, though they do less to address time pressures, which our results suggest are often the more binding margin for mothers. In that sense, policies that lower the time

cost of children may be at least as important as policies that raise household income.

All this relates to a broader structural literature that links childcare technologies, opportunity costs, and bargaining to intra-household specialization (Blundell, Chiappori, and Meghir, 2005; Browning and Gørtz, 2012), partner-specific career dynamics (Adda, Dustmann, and Stevens, 2017; Lise and Yamada, 2019), and the distribution of costs in environments with limited commitment (Voena, 2015; Chiappori and Mazzocco, 2017). The relevance of limited commitment is especially clear in the fertility context. A child changes future time demands, labor-market trajectories, and the value of remaining together, while the associated arrangements regarding childcare and specialization may not be fully enforceable *ex ante*. In that respect, our results should be read as reduced-form evidence on the uneven incidence of the costs associated with children, in a setting where the sustainability of future arrangements is likely to matter.

Several limitations should be noted. Our analysis relies on cross-sectional data and therefore does not correct for selection into fertility: unobserved traits may jointly determine fertility choices and intra-household allocations. Likewise, education, marriage, and fertility are jointly chosen, whereas we treat them as given. The differences we document across fertility groups should therefore be interpreted as structural associations rather than causal effects of childbirth. A second limitation concerns measurement. The PSID lacks detailed information on child-specific inputs for the full sample, requiring us to absorb these into public consumption and home production.

These limitations point toward a broader research agenda. A unified, preferably dynamic, structural framework that jointly models partner search, fertility, labor supply, and human capital accumulation over the life cycle would allow one to go beyond the cross-sectional associations documented here. Such a framework could build on existing work on specialization, career dynamics, and limited-commitment environments cited above, and would provide a natural setting for studying how policy changes affect not only household-level outcomes but also the intra-household distribution of the costs of children. Extending the RP framework to data with richer child-specific expenditure and time records, and to dynamic settings with evolving outside options and marital transitions, is another promising avenue for future work.

In sum, combining marriage-market stability constraints with a collective model of household behavior allows us to uncover the distributional consequences associated with fertility at the individual level. These consequences fall disproportionately on mothers, generating large material and especially time costs. Understanding these asymmetric welfare effects is important not only for interpreting fertility behavior, but also for designing policies that reduce the burdens associated with childrearing and with the unequal allocation of parental time and resources within households.

References

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

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

Becker, G. S. (1974). A theory of marriage: Part ii. *Journal of Political Economy 82*(2), S11–S26.

Blundell, R., P.-A. Chiappori, and C. Meghir (2005). Collective labor supply with children. *Journal of political Economy 113*(6), 1277–1306.

Browning, M., L. Cherchye, T. Demuynck, B. De Rock, and F. Vermeulen (2024). Spouses with benefits: on match quality and consumption inside households. *FEB Research Report Department of Economics*.

Browning, M. and M. Gørtz (2012). Spending time and money within the household. *Scandinavian Journal of Economics 114*(3), 681–704.

Calvi, R. (2020). Why are older women missing in india? the age profile of bargaining power and poverty. *Journal of Political Economy 128*(7), 2453–2501.

Cherchye, L., B. De Rock, and K. Surana (2026). Labor market participation, marriage and individual welfare. *Journal of Political Economy: microeconomics*.

Cherchye, L., B. De Rock, and F. Vermeulen (2012). Married with children: A collective labor supply model with detailed time use and intrahousehold expenditure information. *American Economic Review 102*(7), 3377–3405.

Cherchye, L., B. De Rock, and F. Vermeulen (2026). Individual-level inequality. ECARES Working Paper 2026-10, ECARES, Université libre de Bruxelles.

Cherchye, L., B. De Rock, F. Vermeulen, and P. E. Gobbi (2026). The value of marriage and fertility: A blueprint for a structural approach. ECARES Working Paper 2026-11, ECARES, Université libre de Bruxelles.

Cherchye, L., T. Demuynck, B. De Rock, and F. Vermeulen (2017). Household consumption when the marriage is stable. *American Economic Review 107*(6), 1507–1534.

Chiappori, P.-A. (1988). Rational household labor supply. *Econometrica 56*(1), 63–90.

Chiappori, P.-A. (1992). Collective labor supply and welfare. *Journal of Political Economy 100*(3), 437–467.

Chiappori, P.-A. and M. Mazzocco (2017). Static and intertemporal household decisions. *Journal of Economic Literature 55*(3), 985–1045.

Cortés, P. and J. Pan (2023, 1 December). Children and the remaining gender gaps in the labor market. *Journal of Economic Literature 61*(4), 1359–1409.

Doepke, M., A. Hannusch, F. Kindermann, and M. Tertilt (2023). The economics of fertility: A new era. In *Handbook of the Economics of the Family*, Volume 1, pp. 151–254. Elsevier.

Dunbar, G. R., A. Lewbel, and K. Pendakur (2013). Children's resources in collective households: Identification, estimation, and an application to child poverty in malawi. *American Economic Review 103*(1), 438–471.

Gobbi, P. E., A. Hannusch, and P. Rossi (2026). Family institutions and the global fertility transition. *Journal of Economic Perspectives 40*(1), 47–70.

Goldin, C. (2021). *Career and family: Women's century-long journey toward equity*. Princeton University Press.

Goldin, C. (2024). Nobel lecture: An evolving economic force. *American Economic Review 114*(6), 1515–1539.

Goldin, C. (2025). Babies and the macroeconomy. *Economica 92*(367), 675–700.

Guryan, J., E. Hurst, and M. Kearney (2008, 1 July). Parental education and parental time with children. *The Journal of Economic Perspectives: A Journal of the American Economic Association 22*(3), 23–46.

Kleven, H., C. Landais, and G. Leite-Mariante (2025, 5 September). The child penalty atlas. *The Review of Economic Studies 92*(5), 3174–3207.

Kleven, H., C. Landais, and J. E. Søgaard (2019, October). Children and gender inequality: Evidence from denmark. *American economic journal. Applied economics 11*(4), 181–209.

Lise, J. and S. Seitz (2011). Consumption inequality and intra-household allocations. *Review of Economic Studies 78*(1), 328–355.

Lise, J. and K. Yamada (2019). Household sharing and commitment: Evidence from panel data on individual expenditures and time use. *Review of Economic Studies 86*(5), 2184–2219.

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

Williams, J. R., Y. J. Masuda, and H. Tallis (2016). A measure whose time has come: Formalizing time poverty. *Social Indicators Research 128*(1), 265–283.

A State-level marriage-market sizes

Table 9 reports the number of observations in each of the 26 state-level marriage markets used in the analysis. As explained in Section 3.1, states with fewer than 50 observations are grouped with adjacent or demographically similar states to avoid unrealistically thin marriage markets.

Table 9: Marriage markets

States	N
Alabama	71
Arizona, Colorado, Nevada, New Mexico, Utah	159
Arkansas, Oklahoma	68
California	272
Connecticut, Maine, Massachusetts, New Hampshire, Rhode Island, Vermont	83
Delaware, District of Columbia, Maryland	153
Florida	130
Georgia	130
Idaho, Montana, Oregon, Wyoming	61
Illinois	98
Indiana	109
Iowa	50
Kansas, Nebraska	59
Kentucky, Tennessee	106
Louisiana, Mississippi	155
Michigan	163
Minnesota, North Dakota, South Dakota, Wisconsin	89
Missouri	94
New Jersey, New York	142
North Carolina	185
Ohio	136
Pennsylvania	108
South Carolina	149
Texas	224
Virginia, West Virginia	90
Washington	56
Total	3,127

B Regression results for singles

Table 10 reports the regression results for singles. These estimates serve as a useful benchmark for the main analysis of married couples. They show how fertility and education correlate with consumption and leisure when there is no intra-household allocation mechanism through which the costs of children can be shared.

Table 10: Regression results; singles

	Males		Females	
	Log Consumption	Leisure	Log Consumption	Leisure
High education	177.5***	2.133	132.9***	-0.834
	(31.52)	(1.639)	(51.41)	(1.680)
One child	-89.01	-17.69***	-186.3***	-19.63***
	(89.76)	(5.771)	(51.93)	(3.269)
Two+ children	-179.3***	-32.78***	-303.2***	-19.29***
	(68.68)	(7.775)	(41.08)	(2.717)
High education × One child	-133.2	-4.212	-38.96	-0.191
	(93.06)	(6.108)	(72.77)	(3.431)
High education × Two+ children	-262.6***	18.65**	-129.8*	-0.588
	(78.47)	(9.343)	(67.68)	(3.170)
Age of youngest child	-2.257	0.964**	-1.370	0.495**
	(6.443)	(0.479)	(3.318)	(0.215)
Employed	248.1***	-41.10***	221.1***	-36.16***
	(35.54)	(2.113)	(26.76)	(1.803)
Age	2.915	-0.207*	4.943	-0.0267
	(2.466)	(0.112)	(4.291)	(0.0975)
Home owner	164.3***	-3.276*	174.2***	-4.784***
	(37.84)	(1.710)	(56.31)	(1.481)
Marriage Market FE	✓	✓	✓	✓
Observations	762	762	1076	1076

Notes: Robust standard errors are reported in parentheses. Low education is defined as $\leq$ 12 years of schooling, while high education is defined as $>$ 12 years. The reference group consists of individuals with low education, no children living in the household, and who do not own their home. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

C Material and time poverty, including singles

Figures 6 and 7 complement the main-text distributional analysis by pooling singles and married individuals. They are useful for assessing to what extent fertility-related gradients in consumption and leisure are specific to marriage and to what extent they also appear outside marriage.

Figure 6: CDF of log consumption by number of children

Male

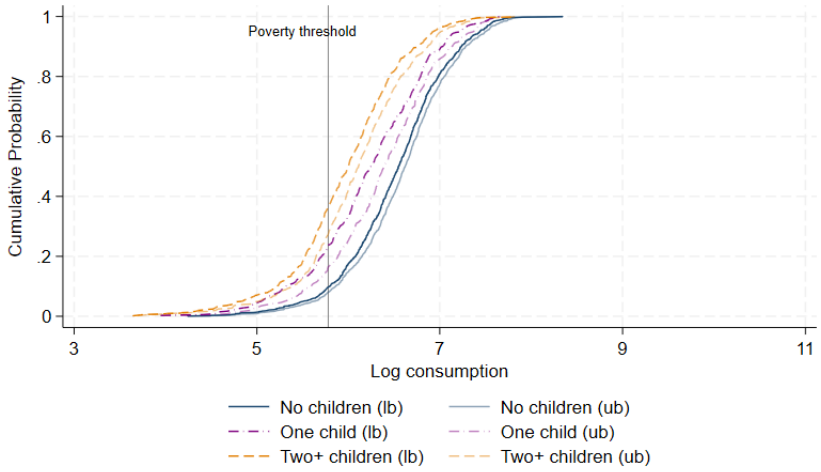

Female

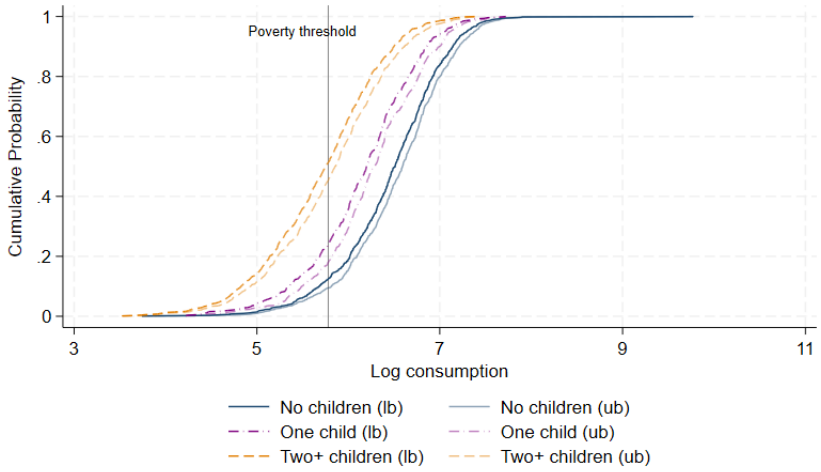

Notes: The material-poverty threshold is set at 60% of the median equivalized consumption in the sample. For single individuals, consumption is directly observed and point-identified. For married individuals, whose consumption is set-identified, we use the midpoint of each person's identified consumption interval when constructing the poverty threshold.

Figure 7: CDF of leisure by number of children

Male

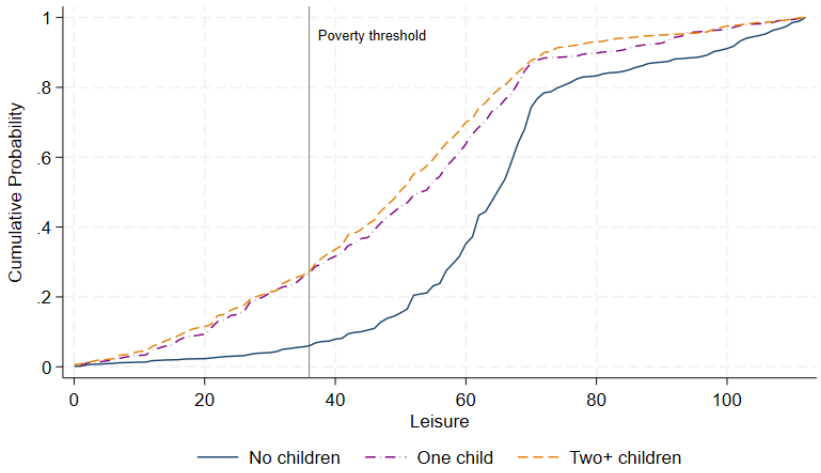

Female

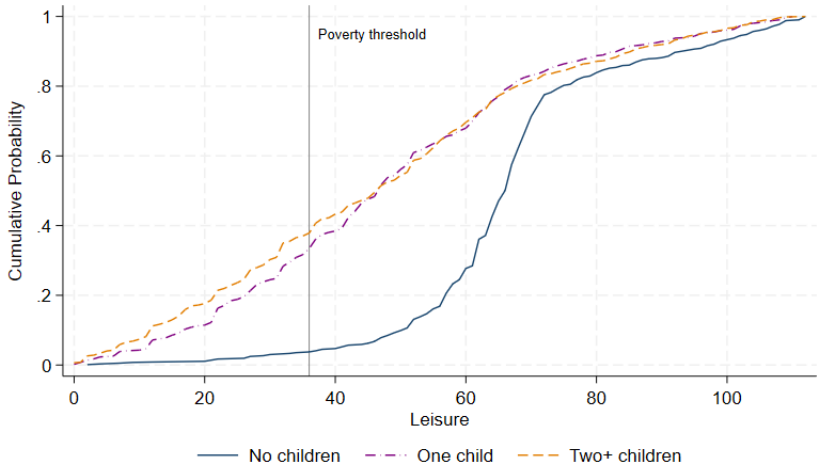

Notes: The time-poverty threshold is defined as 60% of median individual leisure in the sample.