The Child Cost Index*

Christina Patterson
Chicago Booth, CEPR & NBER

Heather Sarsons
Chicago Booth, CEPR & NBER

October 2026

Abstract

Rising child-rearing costs are often cited in accounts of declining U.S. fertility, but there is little evidence on how the prices of child-rearing inputs have evolved in the aggregate. We construct a Child Cost Index (CCI) for 1990–2023 that combines purchased goods and services with parental time and tracks the price of this composite child-rearing bundle relative to adult consumption. Despite rapid inflation in childcare and tuition prices, the goods-and-services CCI closely tracks adult prices, as these increases are offset by children's lower exposure to shelter and by slower price growth elsewhere in the child basket. Households devote substantially more real resources to children than they did in 1990, but this increase parallels the growth of real adult consumption rather than reflecting a child-specific rise in prices. A similar offsetting pattern appears for parental time: the rising value of women's time is largely counterbalanced by declining motherhood penalties in earnings and work hours, so forgone market work adds little to relative price growth. The main departure from this stability comes from the growing amount of leisure displaced by childcare; valuing that time raises the CCI, although the magnitude depends on the shadow value assigned to leisure.

*We thank Violet Hamlin, Diana Pinilla, and Yulia Valdivia Rivera for excellent research assistance. We also thank Otis Reid for helpful comments and discussions. All errors are our own.

1 Introduction

The U.S. total fertility rate has fallen sharply over time and now sits at record lows. The cost of raising children figures prominently in both public discussion and recent economic research as a contributor to the decline (New York Times/Morning Consult, 2018; Pew Research Center, 2024; OECD, 2024). A growing microeconomic literature studies how fertility responds to particular cost margins, including childcare prices, housing costs, and the labor-market consequences of parenthood (Kleven et al., 2019; Couillard, 2025; Dow, 2026; van Wijk and Feijten, 2025). Yet children require a bundle of goods, services, and parental time, and the prices of these inputs need not move together. Rapid growth in the price of one salient input does not, by itself, reveal whether raising children has become more expensive relative to the consumption households would otherwise purchase. Answering that question requires tracking the full set of child-rearing input prices.

This paper constructs a Child Cost Index (CCI) that tracks the price of producing a reference child-year from 1990 to 2023, and compares it with an Adult Consumption Index (ACI) that measures the price of the consumption that child-related spending displaces. The construction follows the logic of consumer price measurement: we aggregate changes in the prices of child-rearing inputs using their expenditure and time-cost shares. The CCI is a macro accounting object, designed to be used alongside, rather than in place of, margin-specific causal estimates. Our central finding is that the monetary price of a child-year has risen at almost exactly the pace of the adult consumption it displaces. Children are expensive, but by the standard tools economists use to measure price change over time, there is little evidence that their relative price has risen dramatically. Only when we price parents' forgone leisure does the relative price of children increase meaningfully.

The empirical construction has two components. For goods and services, we use the Consumer Expenditure Survey to recover the spending basket associated with an age-standardized child-year, allocating household expenditures to children through category-specific rules. We map these expenditures to detailed Consumer Price Index series in a chained Törnqvist index that allows expenditure shares to change over time. The adult benchmark is the marginal spending bundle of married couples without children, estimated from category-specific Engel curves, which approximates how a dollar not devoted to children would otherwise be spent. For parental time, we combine time-use surveys, which measure how many childcare hours parents supply and whether those hours come from market work or leisure, with estimates of the labor-market consequences of motherhood from the Current Population Survey and American Community Survey. We use an approach similar to Kleven (2025), where we compare women and men with children to observably similar women and men without children.

We estimate the impact of children on work and leisure hours at each child age. This allows us to price a forgone work hour by the total earnings loss associated with it, including changes in job characteristics, wage growth, and career progression. Because leisure has no directly observed market price, we report a range of indices that differ in how they value childcare time drawn from leisure.

We find four main patterns. First, the goods-and-services CCI and the ACI both more than double between 1990 and 2023 and track one another closely throughout. Their ratio fluctuates within a narrow range and ends the sample slightly below its 1990 value. This stability survives a battery of perturbations—alternative allocations of housing, alternative adult benchmarks, fixed rather than chained weights, and different weights on the ages of childhood. Housing allocation is the most consequential of these choices, and even it moves the 2023 relative price by only a few percentage points. Relative goods-and-services CCI inflation is somewhat higher among more educated, urban, and older parents, but the differences are small in aggregate terms.

Second, this stability masks large offsetting movements across expenditure categories. Childcare and school-tuition prices rose faster than the adult benchmark and make the largest positive contributions to relative child-price inflation; medical care also contributes positively under our baseline benchmark. Shelter is the most important offset. In our baseline construction, children are assigned the incremental housing cost of an additional bedroom rather than a proportional share of all housing services. Children consequently carry a smaller shelter weight than the adult benchmark, so rapid shelter inflation raises the ACI more than the CCI. Categories in which children have high expenditure shares but more moderate price growth, such as food at home, provide further offsets. The stability of the aggregate index is therefore not evidence that childcare or tuition inflation was small. It is an accounting result arising from differences in expenditure exposure across the full child and adult baskets.

Third, real spending on children increased substantially, but not disproportionately. Real spending per child, deflated by the goods-and-services CCI, rose by roughly one-fifth between 1990 and 2023—and real spending per adult rose by a similar amount. The share of household spending allocated to children is also remarkably stable within family size. Child-rearing has become more resource intensive without becoming substantially more expensive relative to the household's consumption alternative. This distinction between prices and chosen quantities is important for interpreting the rise in observed expenditures on children.

Fourth, the two components of parental time point in different directions. Adding mothers' forgone market earnings raises the level of the price of a child in every year but adds little to the growth of the index. The motherhood penalties in earnings and work hours

declined over the period, so women now give up fewer market-work hours when they have young children, and the effective price of a forgone hour—total earnings losses divided by hours lost—grew considerably more slowly than the ordinary female wage. The share of goods-plus-work child costs attributable to forgone work falls from 33 percent in 1990 to roughly 13 percent in 2023, and the goods-plus-work index remains close to, and for much of the sample slightly below, the goods-only index. Forgone leisure is the more consequential margin. Childcare time increased for both mothers and fathers, and the additional hours came primarily from leisure rather than market work. Leaving forgone leisure unpriced produces an index close to the monetary benchmark; valuing it at the wage raises the growth of the CCI relative to the ACI; and allowing the value of leisure to rise as leisure goods become cheaper relative to other consumption raises it further. Because the data identify the quantity of leisure displaced by children but not its shadow value, we interpret the leisure specifications as a range of possible valuations rather than as alternative estimates of a single observed price.

These results discipline, rather than dismiss, cost-based accounts of declining fertility. They provide little support for a simple narrative in which the average relative price of a fixed-standard child-year rose dramatically. They remain consistent with several mechanisms that an aggregate price index is not designed to capture. Child-related expenses are concentrated in the early years of childhood, when households have lower earnings and limited access to credit. Large, discrete payments for daycare, tuition, or medical care may be more salient than diffuse offsets elsewhere in the household budget. The bundle parents regard as necessary may expand over time, even if the price of a reference bundle does not. And average price movements may mask substantial burdens for households that rely on formal childcare or experience other constraints. The macro evidence thus shifts attention from broad-based average price growth toward the timing, concentration, heterogeneity, and perception of child-related costs.

The paper contributes first to the literature on fertility trends. Classic models emphasize women's opportunity costs and the quantity–quality tradeoff as central determinants of fertility (Becker, 1960; Mincer, 1963; Greenwood et al., 2005; Galor, 2010). Empirical work indeed shows that the timing of fertility responds to women's labor-market opportunities, and was further influenced by medical innovations like birth control that allowed women to take advantage of those opportunities (Goldin and Katz, 2002; Bailey, 2006; Goldin, 2021; Miller, 2011). More recent work argues that these conventional channels cannot fully account for the decline in U.S. fertility since the mid-2000s (Kearney et al., 2022) and that the value of time accounts for part of the decline (Fernández et al., 2026), while other studies estimate fertility responses to individual child-related costs (Couillard, 2025). We complement these

studies by constructing a macro time series for the broader child-rearing price environment. The causal effect of a particular price change and the aggregate evolution of the full set of input prices are distinct objects, and both are relevant for assessing the role of economic costs in fertility trends.

Second, the paper contributes to the literature on household-specific inflation. Existing work constructs price indices for households that differ in income, age, and expenditure composition and shows that these differences matter for welfare and inequality (Hobijn et al., 2009; Jaravel, 2021). We apply the same logic to a demographic object: the child and adult baskets differ systematically in their exposure to individual prices, making headline inflation an incomplete measure of the relative price relevant for child-rearing. We also extend conventional expenditure-based measurement by incorporating nonmarket time. The resulting series provides a macro input for quantitative work on fertility, household production, and demographic change.

Finally, the paper relates to the literature on equivalence scales and the cost of children (Rothbarth, 1943; Deaton and Zaidi, 2002). That literature asks how much additional expenditure a household requires to maintain a given level of welfare after the arrival of a child. Our object is different: we measure how the price of producing a reference child-year evolves over time, not the level of compensation associated with adding a child at a point in time. This distinction allows us to separate changes in input prices from changes in the resources households choose to devote to children.

The remainder of the paper is organized as follows. Section 2 presents the conceptual framework. Section 3 describes the data. Section 4 constructs the goods-and-services CCI. Section 5 incorporates parental time. Section 6 concludes.

2 Conceptual Framework

We begin with a simple conceptual framework that motivates our construction of the Child Cost Index. We adapt the household-production framework of Becker (1965). In the Becker (1965) model, individuals combine goods and time to produce "commodities", which enter utility. We extend this framework so that individuals can produce children (which we refer to as "child services"), also requiring goods and time. The framework serves two purposes. First, it defines the price object we measure: the price of producing a child of fixed quality as input prices change. Second, it clarifies why the relevant empirical object is this child price relative to the price of adult consumption.

The household has preferences $U(c_t, Z_t)$ over adult consumption c_t and child services Z_t, with U increasing in both arguments; the full specification is given in Appendix D. Child

services combine the number of children n_t and the quality, or resource intensity, of each child q_t:

$$Z_t = n_t q_t^{\eta}, \qquad \eta \in (0,1). \tag{1}$$

Child services are increasing in both the number of children and the resources devoted to each child. We assume that children are time-intensive: child quality is produced using both parental time e_t and child-specific market goods x_t,

$$q_t = A_t e_t^{\alpha} x_t^{1-\alpha}, \qquad \alpha \in (0,1). \tag{2}$$

Each child also requires a fixed amount of care time h, independent of quality. A household with time endowment $\overline{T}$, non-labor income V_t, wage w_t, price of adult consumption P_t^a, and price of child market goods p_t faces the full-income budget constraint

$$P_t^a c_t + n_t \left[p_t x_t + w_t (h + e_t) \right] = w_t \overline{T} + V_t \equiv I_t. \tag{3}$$

The term in brackets is the resource price of one child. It consists of market spending on child goods, $p_t x_t$, fixed care time, $w_t h$, and quality-producing parental time, $w_t e_t$.

For a given reference quality $\bar{q}$, the household's minimum price of producing one child of that quality is

$$\mathcal{C}_t(\bar{q}) = \min_{e,x} \left\{ w_t(h+e) + p_t x : A e^{\alpha} x^{1-\alpha} \geq \bar{q} \right\}. \tag{4}$$

We define the Child Cost Index as the price index associated with this fixed-quality cost function:

$$\mathrm{CCI}_t \equiv \mathcal{C}_t(\bar{q}).$$

The minimization problem in equation (4) has the standard Cobb–Douglas dual cost function. At the cost-minimizing input mix, the household equates the ratio of input prices to the ratio of marginal products, so

$$\frac{e_t^*(\bar{q})}{x_t^*(\bar{q})} = \frac{\alpha}{1-\alpha} \frac{p_t}{w_t}.$$

Substituting this input ratio back into the quality-production constraint gives the minimized variable cost of producing one unit of child quality:

$$\pi_t^q = B(A) w_t^{\alpha} p_t^{1-\alpha}, \qquad B(A) = A^{-1} \alpha^{-\alpha} (1-\alpha)^{-(1-\alpha)}. \tag{5}$$

The term π_t^q is the unit cost of child quality: it is the minimum cost of producing one unit of q, before adding the fixed care-time requirement h. Thus the total cost of a child of reference

quality $\bar{q}$ is

$$\mathcal{C}_t(\bar{q}) = w_t h + \pi_t^q \bar{q}. \tag{6}$$

Holding the reference quality and production technology fixed, changes in $\mathcal{C}_t(\bar{q})$ are therefore driven by changes in the prices of child goods and parental time. This is the price variation the CCI is designed to measure.

By Shephard's lemma, differentiating the fixed-quality cost function with respect to input prices gives the growth rate of the index:

$$d \ln \mathrm{CCI}_t = s_t^G \, d \ln p_t + s_t^T \, d \ln w_t, \tag{7}$$

where

$$s_t^G = \frac{p_t x_t^*(\bar{q})}{\mathcal{C}_t(\bar{q})}, \qquad s_t^T = \frac{w_t[h + e_t^*(\bar{q})]}{\mathcal{C}_t(\bar{q})}, \qquad s_t^G + s_t^T = 1.$$

Thus child-cost inflation is a share-weighted average of inflation in child market inputs and inflation in the opportunity cost of parental time.

Equation (7) provides the foundation for the child cost index that we measure in the data. The empirical index generalizes this expression in two ways. First, child market goods consist of many expenditure categories k, each with its own price series $P_{k,t}$. Second, the expenditure shares are allowed to evolve over time. The goods-and-services component of the index is therefore

$$\Delta \ln \mathrm{CCI}_t^g = \sum_k \bar{s}_{k,t}^g \Delta \ln P_{k,t}, \tag{8}$$

where $\bar{s}_{k,t}^g$ is the two-period average of category k's share in the child goods-and-services basket. Equation (8) is the Törnqvist, or chained, price index.

When we add parental time, the full index becomes

$$\Delta \ln \mathrm{CCI}_t = \bar{s}_t^G \Delta \ln \mathrm{CCI}_t^g + \bar{s}_t^T \Delta \ln W_t, \tag{9}$$

where W_t is the opportunity cost of the hours parents devote to children. How we measure W_t in practice—including how the answer depends on whether those hours are drawn from work or from leisure—is discussed in Section 5.

This construction is deliberately a price index rather than a measure of realized spending per child. Realized spending reflects both prices and household choices over the quality of children. If childcare prices rise, households may spend more to maintain the same quality, or they may hold spending fixed and reduce some dimension of quality. The CCI holds quality fixed and asks how the price of delivering that reference child-year changes as input prices change. This is the same logic that underlies the CPI: the object is the price of a fixed

standard.

The framework also explains why we focus on the relative price of children. A fertility decision trades off child services against adult consumption. If the price of children and the price of adult consumption rise at the same rate, children have not become more expensive in the relevant economic sense: the same amount of adult consumption must be forgone to finance a child of fixed quality. What matters is therefore the price of the child bundle in units of the adult consumption bundle. At the level of the model, we define this relative price as

$$\Pi_t = \frac{\text{CCI}_t}{P_t^a}. \tag{10}$$

In the data, P_t^a is not observed directly: we cannot observe the consumption basket parents would have purchased in the absence of children. We therefore construct an empirical counterpart, the Adult Consumption Index, ACI_t, using the consumption bundle of married adults without children as a proxy for the adult consumption that child spending displaces. The empirical relative price that we measure throughout the paper is

$$\widehat{\Pi}_t = \frac{\text{CCI}_t}{\text{ACI}_t}. \tag{11}$$

This object measures how expensive a fixed-quality child has become in units of adult consumption.

Appendix D formalizes this intuition in a CES version of the model. After normalizing the household budget by the price of adult consumption, choices over adult consumption and child services depend on child input prices through the relative price CCI_t/P_t^a. The sign and magnitude of the fertility response depend on preferences and income effects: when child services and adult consumption are substitutes, an increase in the relative price of children lowers the desired allocation to child services. Our goal is not to show the effect of prices on fertility, but rather to measure how the relative price that enters the fertility decision has moved.

The empirical implementation has two main components: the goods-and-services price index, which requires recovering child-specific expenditure bundles and mapping them to observed price series, and the time-cost component, which requires measuring parental hours and pricing them. Section 3 describes the data underlying both; Section 4 constructs and presents results for the goods component, and Section 5 does the same for parental time.

3 Data

We draw on three main datasets. We use the Consumer Expenditure Survey to measure the goods-and-services component of child prices, the Current Population Survey and American Community Survey to price parental time through its effect on earnings and hours, and the American Heritage Time Use Study to measure how parents allocate time. This section introduces each dataset; the measurement choices that turn these data into index components are described in Sections 4 and 5.

Child Spending Shares: Consumer Expenditure Survey (CEX). The Consumer Expenditure Survey is administered by the U.S. Census Bureau on behalf of the Bureau of Labor Statistics (BLS). The survey collects data on household spending, reporting spending in both granular and broad categories. In constructing the CPI, the BLS uses exactly these data to calculate the share of household spending that goes to each category. We use the CEX to recover the share of spending on a child that falls in each category, which corresponds to the weights s_G in equation 8, and the within-goods category shares in the index. We restrict the sample to households with children under the age of 18 and in which all parents are between 20 and 60.[1] Following Lino et al. (2017), we annualize the CEX data, multiplying each subcategory by four (and the rental equivalence of housing, which is measured monthly, by twelve). For comparison, we use the spending of married couples without children, restricting to couples where both partners are between 20 and 60 so that the sample is comparable to the sample of families with kids.

Labor Market Costs: Current Population Survey (CPS) and American Community Survey (ACS). We use the Current Population Survey and American Community Survey, both repeated cross-sections, to measure how a first child affects parents' earnings and work hours. Because we cannot follow the same individuals over time in the CPS and ACS, we compare individuals with young children to observationally similar individuals who do not yet have a child. We use these data to estimate the total earnings lost when having children, which informs the time share s_T, and to estimate the price of each hour of that time, which we use to value the parental hours devoted to childcare.

Leisure Costs: Harmonized Time Use Surveys. The time that parents devote to childcare comes only partly from forgone market work; the remainder comes from leisure

[1] We also exclude households with children 18 or older, as it is unclear whether to treat those children in the same way as younger dependents.

and other uses of time. We document the time costs of children using time-use data. Time-use data have been collected systematically since 2003 in the American Time Use Survey (ATUS), conducted by the Bureau of Labor Statistics. We make use of two additional surveys that provide time-use data going back to the early 1990s: the Americans' Use of Time Survey and the Time Use in Social and Economic Accounts.[2] Each survey targeted a nationally representative sample, interviewed one individual per household (typically an adult, 18+), and used a 24-hour recall of activities.

The American Heritage Time Use Study (AHTUS) harmonizes activities across these surveys so that we can compare time spent in different activities across years. All surveys measure "primary" childcare—activities for which childcare is the primary activity—but only the ATUS measured "secondary childcare"—activities for which a child is present but which are not the primary activity.[3] We therefore measure the gap between primary and secondary childcare in the ATUS and extend the estimates backward to have estimates of secondary childcare throughout.

4 The Price of a Child's Expenditure Bundle

4.1 Methodology: Inferring Spending on Children

We begin by looking at the raw differences in spending between families with children and families without children. As Figure 1 illustrates, the composition of household spending shifts markedly with the presence of children: relative to childless married couples, households with children allocate a larger share of spending to food and education and childcare, and a smaller share to housing and recreation. These differences reflect a mix of spending done directly for children, changes in adults' own consumption in response to having children, and selection into parenthood. To isolate what households spend on children specifically, we devise allocation rules that apportion household expenditures across family members.[4]

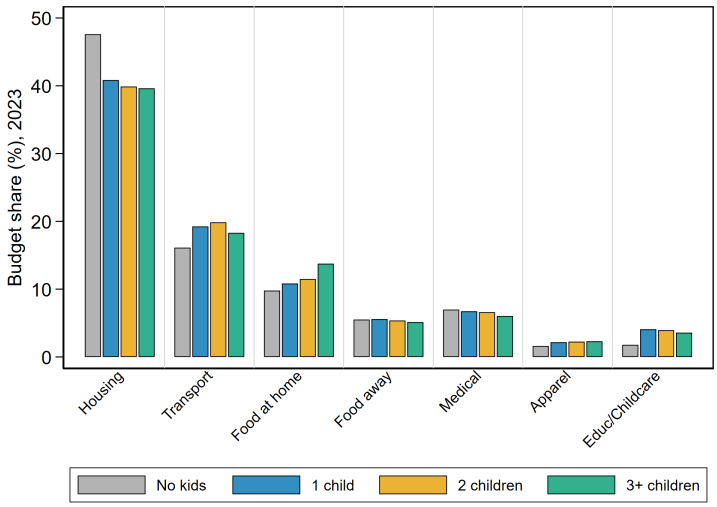
Some spending, such as clothing, is reported separately for adults and children in the

[2]An additional time-use survey is the National Human Activity Pattern Survey. We cannot use this survey because it does not record respondents' marital status or number of children.

[3]Examples of primary childcare are reading a child a story or helping a child get dressed. An example of secondary childcare is reading while monitoring children playing.

[4]Appendix Figure E.5 reports a descriptive pseudo-event study of household spending around the birth of a first child. Total spending rises after childbirth, and the increase is concentrated in intuitive categories: food, education and childcare, healthcare, and apparel. Housing and transportation also move, though these responses are harder to interpret as direct child expenditures because they reflect broader household adjustments. We therefore do not use these event-study coefficients to construct the child basket. Instead, the figure serves as descriptive evidence that the arrival of children is associated with meaningful and intuitive changes in household spending.

Figure 1: Budget shares differ for families with and without kids



Notes: Bars show category-level budget shares in 2023, computed from the Consumer Expenditure Survey (CEX). Within each spending category the four bars correspond to households grouped by the number of resident children: married couples without children (gray); and married couples with one child (blue), two children (gold), and three or more children (green).

CEX, while most spending, such as food, is not. We therefore devise allocation rules to apportion household spending across members of the household, adapting the rules used by Lino et al. (2017) in the USDA's "Expenditures on Children by Families."[5] We explain how we allocate spending in each broad category below:

Food To determine what spending was for the children, we follow the United States Department of Agriculture (USDA)'s allocation of food for children by age and number of children (Lino et al., 2017).[6] When developing their own price of a child, Lino et al. (2017) calculated what percentage of a family's spending on food should go to each child based on the age of the youngest child and how many children are in the family, derived from the December 2015 USDA Food Plans: Cost of Food at Home Report. [7] We adapt those percentages and apply them to our data by matching them to the number of children in each

[5]In accordance with how the USDA annualized CEX data for their calculation of the cost of having a child, we multiplied each subcategory by four to annualize the data, with the exception of the rental equivalence of housing, which was measured monthly and so was annualized by multiplying by twelve.

[6]We combine all spending for food at home into one subcategory, matching a special aggregation class in the CPI for food at home. Similarly, all alcoholic beverage spending is aggregated into a single subcategory, matching the aggregation in the CPI. The final subcategory is the expenditure class for food away from home.

[7]The USDA releases monthly reports on the cost of four different food plans. The food plans' market baskets were last updated in 2007 (excluding the Thrifty Food Plan, which was updated in 2021), with monthly reports generating the price of the food baskets adjusted for inflation.

family and the age of each child within the family. For example, consider a family with three kids aged 3, 7, and 10. The three-year-old receives 14% of the family's spending on food, the seven year old receives 17% of the family's spending on food, and the ten year old receives 20% of the food spending. All remaining spending on food is allocated to the parents. We adopt the same allocation rules for food at home as we adopt for food away from home. All spending on alcoholic beverages is allocated to the parents.

Clothing Within the CEX, children's clothing is reported separately from adults' clothing. We allocate it proportionally to all children under the age of 18 within the household. Jewelry and watches are allocated proportionally across the entire family.

Transportation To allocate transportation expenses, we follow Lino et al. (2017), assigning shares to each child based on the age of the youngest child and the number of children in the household. In aggregate, this allocates 37.5 percent of total transportation spending to children, divided evenly among them, and the remaining 62.5 percent to parents.[8]

Healthcare To allocate healthcare expenses, we use Lino et al. (2017)'s shares, which are derived from Medical Expenditure Panel Survey (MEPS) data and vary by the age of the youngest child and the number of children in the family.[9] For example, in a family with one child aged 7 and one aged 12, 26 percent of healthcare spending is allocated to the 7-year-old and 19 percent to the 12-year-old.

Education/child care For education and childcare spending, all educational supplies are divided evenly among the children. Nursery school tuition and childcare expenses are divided evenly among children 0-4. In the case of families with positive spending and no children 0-4, the spending is divided evenly among all children. Primary and secondary school tuition is allocated evenly among children aged 5-17. Similarly, in the case of families reporting positive K-12 spending and no children 5-17, the spending is allocated evenly among all

[8] A U.S. Department of Transportation study classifies 75 percent of family transportation as child-rearing related (Santos et al., 2011). Lino et al. (2017) assume that parents and children share equally in child-rearing trips, assigning each group 50 percent of that total on the grounds that such trips are used jointly by family members. This yields 50% × 75% = 37.5% of total transportation to children and 25% + 37.5% = 62.5% to parents.

[9] Due to an issue in the duplication of the same ELI in multiple expenditure classes in the healthcare spending category, we did not use expenditure classes for healthcare. Instead, we used an aggregated healthcare spending variable reported in the FMLI table (the CEX expenditure summary table), which includes both out-of-pocket expenses and insurance payments. The MEPS is a nationally representative survey of the civilian non-institutionalized population that collects detailed information on healthcare utilization and expenditures, conducted by the U.S. Department of Health and Human Services.

children. Since we exclude families with children above 17, all college and technical school tuition is allocated to the adults.

Housing In accordance with how the CPI measures housing inflation, we measure the monthly price of owning a home as the estimated monthly rental price of an owned home, and the price of renting simply as the paid rent. We allocate housing to children in two ways. In our baseline estimates, we assign to children the incremental price of acquiring another bedroom. The key observation underlying this approach is that children do not consume a proportional share of all housing services. They share common spaces—kitchens, living rooms, bathrooms, and the lot—with their parents but require a bedroom of their own. The additional-bedroom approach therefore captures only the incremental expenditure a household incurs by living in a home with one additional bedroom, attributing that cost to the child who necessitates it. We use this as our baseline because it most closely mirrors the incremental housing decision a household faces at the birth of a child: the family must acquire one additional bedroom, not a proportional increase in all housing services.

We estimate the additional cost of a bedroom from the CEX using the regression

$$\log(\text{Housing Spending}_i) = \sum_{k=2}^{5+} \beta^k \, \mathbb{I}(\text{Bed}_i = k) + \gamma_{\text{inc}} + \delta_{\text{edu}} + \lambda_{\text{region}} + \varepsilon_i, \tag{12}$$

where $\mathbb{I}(\text{Bed}_i = k)$ indicates that household i lives in a home with k bedrooms, and the controls are income groups, maximum parent education, and census-region fixed effects.[10] The estimated coefficients give the incremental spending associated with an additional bedroom. Specifically, they provide the difference in log housing expenditures relative to a single-bedroom home due to changes in home size.[11] The estimation is done for all households, including those without any children. The estimation is done separately for rent/rental equivalence, housing insurance, fuel and utilities, and household operations and furnishings, allowing these components of housing to scale differentially with the number of bedrooms.

We also consider results when we divide housing expenses equally across all household members. This per-capita allocation likely overstates the amount of the household's housing expenses that were driven by the presence of the child, since housing is partly a household public good (e.g. kitchens, living areas, and the land are shared regardless of family size). Therefore, dividing total costs by the number of household members attributes to each child

[10] Income is broken out into 50 quantile groups. Parental education attainment is broken out into three groups: less than college, college, and more than college.

[11] The predicted price of bedroom k is $H_k = e^{\beta_0+\beta_k}$, and the additional price of the $(k+1)$th bedroom is $H_{k+1} - H_k$. On average, the price of an additional bedroom has risen from about \$1,000 per year in 1990 to nearly \$4,000 in 2023. See Appendix Figure A.8.

a share of fixed household infrastructure that the parents likely would have consumed in the absence of children. We show how the relative child-price series varies across these three housing-allocation approaches in Appendix Figure A.5.

Miscellaneous other spending The remaining categories are allocated per capita across all members of the household.[12] The one exception is tobacco spending, which is allocated completely to the parents. Note, however, that we find that the results are robust to alternate allocation rules here that tilt the recreation spending within this group towards the children.

Because many of the allocation rules (food, healthcare, transportation) depend on the age of the child, we first compute per-child spending separately by child age and then aggregate across ages, giving each year an even weight. This ensures that changing weights on different categories over time do not reflect changes in the age distribution of children. As fertility has declined, the average age of children under 18 has shifted upward. Since older children require a different set of inputs, this could mechanically alter the aggregate spending basket even if within-age-group expenditure patterns were unchanged. Together, these allocation steps yield the time-varying, category-level expenditure shares $\bar{s}^{g}_{k,t}$ that form the weights in equation (8).

The adult-consumption benchmark. The relative price in equation (11) requires a price index for the adult consumption that is displaced by child spending. The relevant counterfactual is not the entire consumption basket of adults without children, but the adult consumption parents would purchase at the margin if resources were not devoted to children. We therefore construct the baseline Adult Consumption Index (ACI) from the marginal spending bundle of married couples without children.

To estimate this marginal bundle, we use married childless couples aged 25–40 and estimate category-specific budget-share Engel curves:

$$s_{c,it} = \alpha_c + \beta_c \ln(\text{TotExp}_{it}) + f(\text{Age}_{it}) + \mu_{c,r} + \tau_{c,t} + \epsilon_{c,it}, \tag{13}$$

where $s_{c,it}$ is household i's budget share in category c. The coefficient β_c captures how the category's budget share changes with total expenditure. Since spending in category c is $s_{c,it}\text{TotExp}_{it}$, the implied share of an additional dollar allocated to category c is

$$m_{c,t} = \widehat{s}_{c,t} + \widehat{\beta}_c, \tag{14}$$

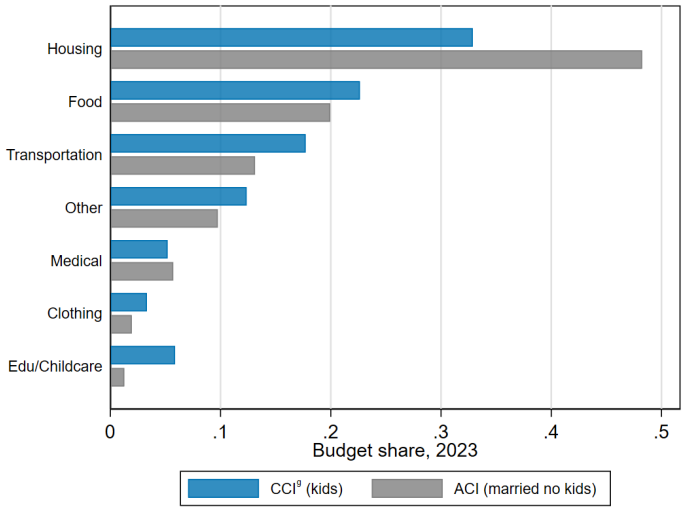
[12]These categories include personal care products, personal care services, miscellaneous personal services, miscellaneous personal goods, video and audio, pets, their products, and their services, sporting goods, other recreation services, photography, and recreational reading.

where $\widehat{s}_{c,t}$ is the predicted budget share for the median married childless couple in year t. We set the small number of negative marginal weights to zero and renormalize the weights to sum to one. These $m_{c,t}$'s are the expenditure weights in our baseline ACI.

As robustness checks, we also construct adult benchmarks using the average spending bundle of married childless couples and using parents' own spending net of the expenditures allocated to children.[13] As shown in Appendix Figure A.5, the resulting relative-price series is very similar across these alternatives.

Figure 2 compares the expenditure shares in the child basket to those in the adult benchmark. Children's weights are tilted toward food and education/childcare and away from housing, whereas the adult bundle loads more heavily on housing.[14] This figure pools all households with children (including single parent households); we separate the weights for single-parent from all households with children in Appendix Figure A.3.

Figure 2: Expenditure shares: weights in the child basket versus the adult benchmark



Notes: The figure compares expenditure weights in two 2023 consumption baskets: the child basket underlying the goods-and-services Child Cost Index (CCI^g, blue) of all household with kids, and the adult benchmark basket underlying the Adult Consumption Index (ACI, gray), where the ACI represents the marginal consumption bundle of married adults without children. The horizontal axis is each category's budget share. Shares are constructed from the CEX using the paper's allocation rules that map household expenditures to children.

[13] Appendix Figure A.4 compares the marginal ACI weights to the average spending shares of married childless couples. The two baskets are broadly similar, but the differences are informative. Relative to the average childless bundle, the Engel-curve marginal bundle places less weight on food at home, food away from home, and medical care, and more weight on shelter, utilities, vehicle insurance, public transportation, apparel, and video/audio goods.

[14] Education for couples with no children represents university tuition.

From expenditures to prices. The final step in constructing the price index is to link the spending categories reported in the CEX (reported at the expenditure-class level) to the categories collected in the CPI. We follow the crosswalk provided in Jaravel (2024) going back to 1999, and for 1990–1998 we manually extend that crosswalk using the data dictionary for the CE Public Use MicroData and the CPI's concordance to match older variables. To account for noise, we pool several years of data and construct shares as rolling averages.[15] Because not all child-relevant CPI subseries are available for the full historical window, we backcast missing early observations using hierarchical proxy series, linking each missing subseries first to a conceptually close parent CPI aggregate and falling back on headline CPI growth otherwise. The goods component of the index is then

$$\Delta \ln P_t^g = \sum_k \bar{s}_k \, \Delta \ln P_{k,t}^g, \tag{15}$$

where $\bar{s}_k$ are the (time-varying, Törnqvist) expenditure shares from the CEX and $P_{k,t}^g$ is the price series for category k. To account for the fact that selection into who has a child may have changed over time, we hold household composition fixed at its 1990 distribution, reweighting each year's households by parent education, age, race, and urban status. The index therefore reflects price and spending changes within cells, not compositional shifts across them.[16]

As an internal validation of the crosswalk from CEX expenditure classes to CPI subseries, we apply our methodology to the full household spending basket rather than to children's spending alone. The resulting aggregate price index closely tracks headline CPI throughout the sample period (see Appendix Figure A.1).

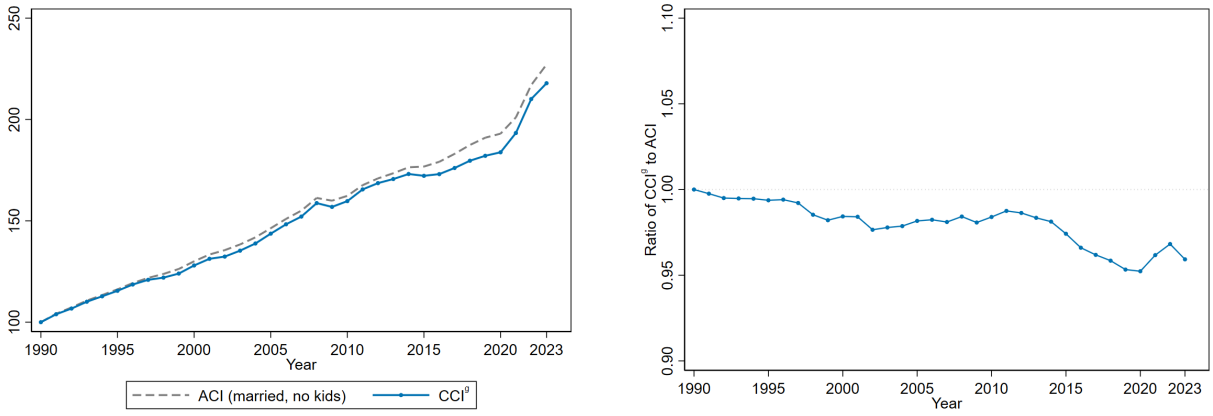
4.2 Results

Three main findings emerge from the goods-and-services component. First, the relative price of the child-rearing bundle has remained essentially stable since 1990, a result that holds across a wide range of specification choices. Second, this stability reflects offsetting forces rather than stable prices: rapid price growth in child-related categories such as childcare and tuition is offset by children's lower exposure to other fast-rising categories, such as housing. Third, while real spending on children has risen, it has done so in step with real

[15] For 1999–2023, we average the four years prior to the given year through the given year; we exclude 1993 and 1994 due to data-quality issues. For 1990-1998, data from 3-5 of the years surrounding the given year are pooled together. Before averaging, we apply CPI adjustments to ensure all data are in the given year's dollars.

[16] In robustness checks, we also re-weight to hold fixed the number of kids, and allow household composition to vary by year.

Figure 3: Price index for child goods and services, CCI_t^g, 1990–2023



(a) Levels

(b) Relative to adult benchmark

Notes: Panel (a) plots the chained goods-and-services Child Cost Index and the Adult Consumption Index, both normalized to 100 in 1990. Panel (b) plots the ratio CCI^g/ACI; values above one indicate that the price of children's goods and services rose faster than the adult benchmark. The ACI is constructed from the marginal consumption bundle of married adults without children.

adult spending, suggesting that child-rearing has become more resource-intensive over time without becoming relatively more expensive.

Result 1: The relative price of children's goods and services is largely stable. Figure 3 presents the central goods-market result. Panel (a) plots the chained CCI^g together with the Adult Consumption Index from 1990 to 2023, with both series normalized to 100 in 1990. The two indices more than double over the sample period and track one another closely throughout. Panel (b) plots the economically relevant object directly: the ratio CCI^g/ACI. This relative price remains close to one and even declines slightly after 2010. By 2023, the price of a child's goods-and-services bundle is about four percentage points below the adult bundle. Thus, in units of adult consumption, a parent in 2023 faces nearly the same market price of children as a parent in 1990.

This conclusion is robust to a range of alternative modeling choices. We begin by probing housing as it is a large component of household spending and is one of the fastest-rising categories in the CPI. We then turn to the construction of the index and to selection into parenthood.

Our baseline allocates the price of an additional bedroom to children. This could understate children's housing costs if couples with children consume a different type of housing that experienced more rapid price growth. We examine two versions of this concern. First, parents with young children may move to areas with better schools that also have higher

house prices. Combining test score data with school-district level house price data, we plot the school-district median home value and gross rent for school districts in the top 10, middle 10, and bottom 10 percent of the test score distribution in Appendix Figure C.1.[17] Although higher-performing school districts have more expensive housing than lower-performing districts on average, the gap has not widened over time. If anything, housing prices in the middle and bottom-performing school districts saw the largest increase in prices over time.

Second, even if a child requires only one additional bedroom, larger homes may have appreciated faster than smaller ones. Using the Zillow Home Value Index by bedroom count for cities and ZIP codes, available from 2000 onward, we construct the premium associated with moving from a $(b-1)$-bedroom unit to a b-bedroom unit, described in more detail in Appendix C.2. Appendix Figure C.2 shows that the premium for a second, third, and fourth bedroom declined between 2000 and 2023.

Appendix Figure A.5 then summarizes a broader set of perturbations to the baseline construction by reporting the 2023 value of CCI^g/ACI under alternative specifications. The solid vertical line marks the baseline estimate (parity between child and adult price growth is at 1). The exercises vary the allocation of housing expenditures to children, the adult counterfactual used to construct the ACI, the expenditure weights used in the price index, the timing weights placed on different child ages, and the observable composition of parents. Under each assumption, the CCI remains within 6-7 percentage points of the ACI.

Replacing the additional-bedroom baseline with an estimate based on differences in housing spending by the number of kids in the household further lowers the 2023 relative price, while allocating housing on a per-capita basis raises the ratio by roughly two percentage points. Thus even a specification that assigns children a relatively high share of housing expenditures does not generate a large divergence between child and adult price growth. The choice of adult benchmark also does not make a substantial difference. Using the average spending bundle of married couples without children produces 2023 ratios close to 1.[18]

The fixed-weight exercises clarify the role of expenditure-weight updating in the relative-price calculation. In these specifications, we hold expenditure weights fixed symmetrically in both the child index and the ACI, either at their 1990 or 2023 values, and then compute the ratio CCI^g/ACI. Holding 1990 weights fixed yields a ratio slightly above one, reflecting that, at 1990 expenditure shares, children were more exposed than adults to categories with high cumulative relative price growth. The lower chained ratio implies that, over time, the child basket shifted away from these categories relative to the adult basket. Holding the 2023 weights fixed yields a relative price close to the chained baseline. The chained index

[17] Appendix C.1 describes this exercise and data in detail.

[18] Appendix Figure A.6 shows how the goods-and-services CCI shifts under each housing measure.

remains our preferred object because it allows both the child and adult baskets to evolve as measured expenditure shares change over time. Similarly, discounting child ages at a 5 percent rate, which places greater weight on prices incurred earlier in childhood, raises the relative price only modestly.

Finally, we probe whether allowing the observable composition of parents and childless couples to vary over time results in a different ratio. In our baseline specification, we hold the composition of parents fixed at its 1990 distribution across age, education, race, and urban/rural status in order to ensure that selection into parenthood does *not* influence the ratio. However, allowing these observable characteristics to adjust again influences the ratio of the CCI to the ACI relatively little.

Taken together, these checks support the conclusion that while reasonable alternatives move the 2023 relative price by a few percentage points, there is no evidence that the monetary price of a child-year has meaningfully pulled away from the adult consumption benchmark.

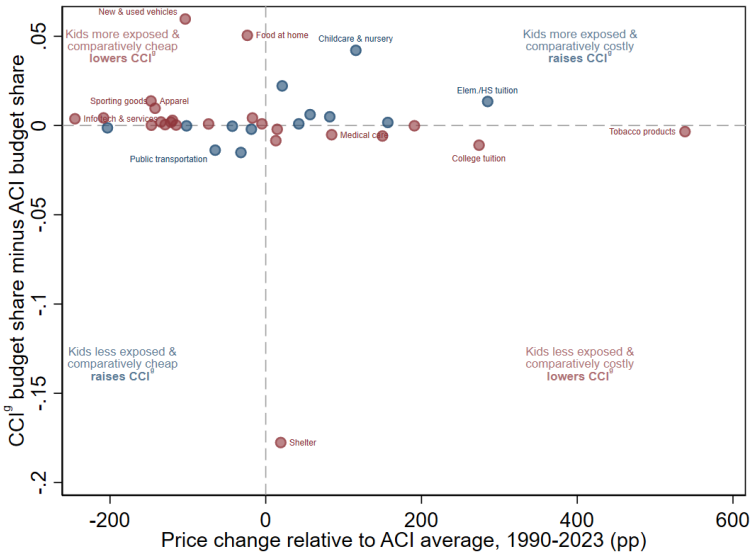
Result 2: Aggregate stability reflects offsetting forces, not uniformly stable prices. The stability of the aggregate index is perhaps surprising in light of the rapid increase in childcare and tuition prices. Appendix Figure A.2 shows that childcare and tuition prices have increased roughly fourfold and sixfold, respectively, since 1990, while the headline CPI index doubled. Figure 4 shows why those salient price increases do not translate into a large increase in the overall relative price of children. We decompose the cumulative child–adult inflation gap into category-level contributions. For category k, the contribution is

$$\text{Contribution}_k = \left(\bar{s}_k^{\text{child}} - \bar{s}_k^{\text{adult}}\right)\left(\Delta\pi_k - \Delta\pi^{\text{ACI}}\right),$$

where $\bar{s}_k^{\text{child}}$ and $\bar{s}_k^{\text{adult}}$ are the average expenditure shares in the child and adult baskets, and $\Delta\pi_k - \Delta\pi^{\text{ACI}}$ is cumulative price growth in category k relative to the adult benchmark. A category raises CCI^g relative to the ACI when children are disproportionately exposed to a category whose price rose faster than the adult benchmark, or when adults are disproportionately exposed to a category whose price rose slowly. A category lowers the relative price of children when the opposite is true. The decomposition is therefore an accounting exercise that separates two forces: differential exposure and differential price growth.

The figure shows that large positive and negative forces offset one another. Education and childcare are the dominant positive contributors: the child basket devotes a meaningful share of expenditures to childcare and school-related services, while the adult alternative consumption bundle devotes very little to these categories. Their prices also rose rapidly over the sample period. Medical care also contributes positively, reflecting the fact that

Figure 4: Category contributions to the child–adult inflation gap



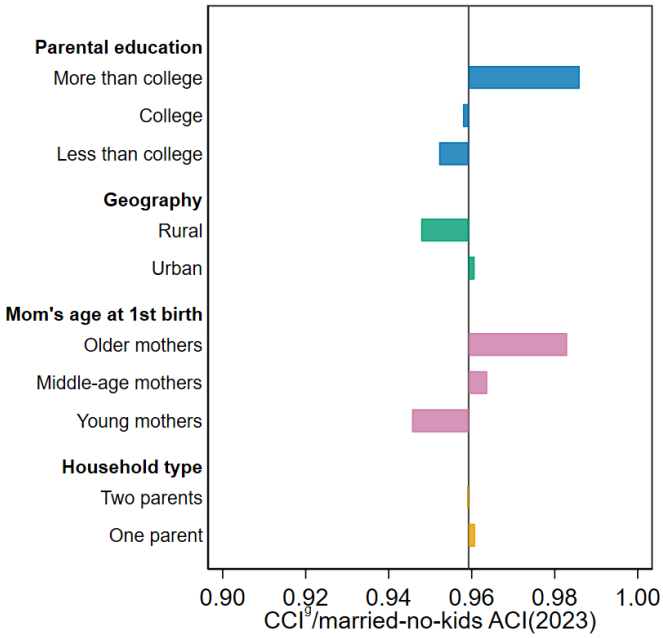
Notes: Each marker corresponds to a spending category. The horizontal axis is cumulative price growth in that category relative to the ACI from 1990 to 2023. The vertical axis is the category's average budget share in the child basket minus its average budget share in the adult benchmark. Marker area is proportional to the category's average budget share in the child basket. Navy markers make positive contributions to CCI^g relative to the ACI; maroon markers make negative contributions. Shares are constructed from the CEX and prices from BLS CPI subindices.

kids have a higher weight than the alternative bundle on this category with relatively fast price growth.[19] These upward pressures are countered by categories in which children have either lower exposure to fast-rising prices or higher exposure to relatively slow-rising prices. Shelter is the most important example. Because our baseline allocates to children only the housing price associated with an additional bedroom, rather than a proportional share of the entire housing bundle, children have lower exposure to housing than the adult benchmark. Since shelter prices rose from 1990 to 2023, this lower exposure raises the ACI more than the child index. Other categories, such as food at home, also pull the index toward parity because children are relatively exposed to categories with more moderate price growth. The near-stability of CCI^g/ACI is therefore not evidence that childcare or tuition inflation was small. It is instead a result of aggregating over fast inflation in highly visible child-related categories and other parts of the child and adult consumption baskets.

An implication of this decomposition is that relative child-price inflation should be higher for groups whose child-rearing baskets place more weight on the categories that became rel-

[19]This category looks different when we use the average spending of married childless households as the alternative consumption bundle. This is because on average, adults have a higher share of their spending on medical expenses than children (even if their marginal spending is not as tilted in this direction). See Appendix Figure E.4.

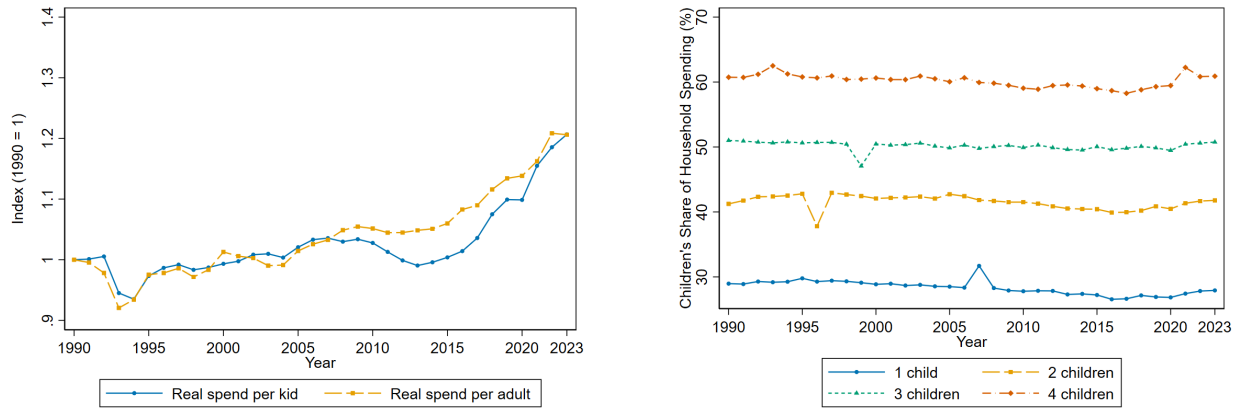
Figure 5: Heterogeneity in relative price growth by observables



Notes: Each bar shows the 2023 relative price of children for a subgroup, measured as the ratio of that subgroup's goods-and-services Child Cost Index to the marginal-bundle Adult Consumption Index. The solid line marks parity between child and adult price growth. Bars extending to the right of the baseline indicate higher relative child-price inflation than the all-children index; bars extending to the left indicate lower. The dashed vertical line is the all-children baseline index. Subgroups are defined by parental education, urban/rural residence, and mother's age at first birth. Indices are constructed from the CEX using the baseline allocation rules.

atively expensive. Figure 5 examines this by constructing subgroup-specific CCI^g indices by parental education, urban status, and mother's age at first birth. The resulting heterogeneity is directionally consistent with the decomposition. Households with more than a college degree face the highest relative price growth among the education groups, and urban households face higher relative child-price inflation than rural households. The index is also higher for older mothers at first birth, a group whose child-rearing basket is more likely to be tilted toward formal childcare, education, and urban services. At the same time, the magnitudes of these differences are small.

Result 3: Real spending on children rose, but proportionally to adult spending. A stable relative price does not mean that child-rearing has remained unchanged. Parents may choose to devote more real resources to each child as living standards rise or as norms of child investment change. In the model, such increases correspond to higher resource intensity—or chosen quality—rather than to a higher price of a fixed-quality child. We therefore ask whether real spending per child, deflated by CCI^g, has risen disproportionately



(a) Real Spending

(b) Household share of spending

Figure 6: Real spending on children and adults, 1990–2023

Notes: Panel (a) plots indices of real spending per child and per adult, both normalized to one in 1990. Real spending per child is nominal child spending deflated by the goods-and-services Child Cost Index (CCI^g); real spending per adult is nominal adult spending deflated by the Adult Consumption Index (ACI). Panel (b) plots the share of total household spending allocated to children over time, separately by family size. All series are constructed from the Consumer Expenditure Survey (CEX) using the baseline allocation rules.

relative to real adult consumption.

Indeed, Figure 6 shows that real spending on children has increased. Panel (a) plots real spending per child, deflated by CCI^g, together with real spending per adult, deflated by the ACI. Both series rise by roughly 20% between 1990 and 2023 and move closely together over the medium run. Thus, households devote more real resources to children today than in 1990, but this increase appears to be part of a broader rise in household consumption rather than a child-specific expansion in required market inputs.

Panel (b) reinforces this interpretation by plotting the share of total household spending allocated to children, separately by family size. Children absorb a large share of household resources in every year of the sample, but that share is strikingly stable over time within family size. Households with one child allocate just under 30 percent of total spending to that child in both 1990 and 2023. The sacrifice of adult consumption associated with children is therefore large, but it is not a recent development. Note that this stability is not mechanically imposed by the allocation procedure: the largest categories are assigned using category-specific rules based on age, household composition, marginal housing prices, or external allocation shares, rather than by imposing a constant child share over time.

Panel (b) also reveals clear economies of scale in child-rearing. Moving from one to two children raises the child share of household spending from roughly 30 percent to roughly 40 percent, far less than a doubling, and the incremental spending share falls further for larger

families. Appendix Figure A.11 unpacks these scale economies by category for children aged 0–5. Holding household and child characteristics fixed, per-child nominal spending is lower for children with older siblings than for otherwise similar first children in nearly every category. The decline is especially large for apparel and education/childcare, and is also visible for housing, transportation, and medical spending; food displays much smaller scale economies. These patterns suggest that many child-related inputs are shared across siblings, reused across children, or purchased less intensively for later-born children.

5 Accounting for Parental Time

Children require not only goods but also parents' time (e_t in our model). Valuing parental time requires two main ingredients: the number of hours parents spend on childcare, and the value of each of those hours. Because childcare time is drawn from two distinct margins—hours reallocated from market work (e_t^M) and hours reallocated from leisure (e_t^L)—we price each margin separately using its own opportunity cost. We first describe how we measure hours and price each margin, and then present the updated child cost index that includes parental time.

5.1 Methodology

We begin by showing descriptively how time allocation between parents and non-parents differs, and then discuss the methodology for pricing those hours.

Measuring childcare hours. We use the time-use surveys to estimate how children alter parents' allocation of time across childcare, market work, leisure, and home production. For individual i, we estimate

$$\text{Hours}_i \;=\; \sum_{b\in\mathcal{B}} \gamma_b\, N_{i,b} \;+\; \psi_2\, \mathbf{1}\{N_i = 2\} \;+\; \psi_3\, \mathbf{1}\{N_i \geq 3\} \;+\; f(\text{Age}_i) + \theta_{\text{educ}} + \lambda_t + \delta_d + \varepsilon_i, \quad (16)$$

separately by gender and time-use category. The outcome is hours per day in the indicated category. $N_{i,b}$ is the number of the respondent's own resident children in age bin $b \in \mathcal{B} = \{0\text{–}5,\ 6\text{–}11,\ 12\text{–}17\}$ and N_i is the total number of children in the household.

The early (pre-2003) survey waves of the AHTUS only ask the age of the youngest child in the household, and whether there are children aged 0-5 and 6-17. We therefore cannot estimate time inputs separately for each child age bin in the 1990s data as we can in the later years. Instead, we recover them in two steps. First, using the pooled 1993-1995 AHTUS

waves, we regress daily hours in each activity on two household-type indicators—youngest child under 5, and youngest child 6–17—along with family-size controls, age, education, and day-of-week fixed effects. These coefficients, shown in Appendix Table B.2, give the average change in parents' time use relative to comparable childless adults, among each household type. However, they do not tell us how that time is divided across children of different ages. The second step is therefore to run both the coarse household-type specification and our main specification (equation 16) using the earliest ATUS waves (2003–2007). For each age bin, we compute the ratio of the per-child coefficient to the corresponding coarse household coefficient. This ratio tells us the share of the total household effect accounted for by one child in each age bin. Multiplying the 1990s AHTUS coarse coefficients by these ratios converts them into per-child time inputs by age bin for the early 1990s. The main assumption we need to make is that although the overall amount of time parents devote to children changed between the early 1990s and the 2000s, the way that time is divided across children of different ages did not. We assess this assumption directly by re-estimating the ratios in later ATUS windows (2008–2015 and 2016–2023) and find them stable over those two decades.[20]

The sample consists of married and cohabitating couples aged 20–60 with spouses aged 20–60, with zero to three children. We control for parent age and education fixed effects, as well as day of week effects. The coefficient γ_b thus measures the daily hours associated with one child in bin b, relative to comparable childless adults. The family-size indicators ψ_2, ψ_3 absorb economies of scale in child-rearing, so that γ_b is identified from variation in the age composition of children holding family size fixed. Under this parameterization, γ_b corresponds to the time input of a first child in bin b, with subsequent children adding $\gamma_b + \psi_2$ (for a second child) or $\gamma_b + \psi_3$ (for a third or higher-order child). The maintained assumption for the index is that, conditional on the controls in equation (16), the differences in daily time use associated with having a child, relative to similar married individuals without children, approximate the childcare-related time reallocation induced by children.

We consider two definitions of childcare time: a *strict* classification, which counts as childcare only time whose main reported activity is childcare-related (e.g. bathing a child), and a *broad* classification, which adds time in which childcare is a simultaneous, supervisory activity (e.g. eating dinner with children).[21]

Panel A of Table 1 reports estimates using only the 2003-2023 ATUS data, where the

[20] Standard errors for the resulting imputed inputs are computed by bootstrap, resampling households within survey years in the AHTUS and ATUS independently.

[21] Because the ATUS records secondary childcare only for household children under age 13, the broad time definition is incomplete for the 12–17 bin. For that bin we scale the strict estimate by the broad-to-strict ratio estimated in the 6–11 bin. In addition, the pre-2003 survey years do not calculate broad childcare. We therefore again scale the estimates using the post-2003 values.

data is most complete. For women, each child under five is associated with 2.18 additional hours of childcare per day, 1 fewer hours of market work, and 1.11 fewer hours of leisure. As children age, the time associated with childcare declines and women experience a smaller change in work and leisure time. For men, the reallocation is smaller and comes almost entirely from leisure: childcare rises by 1.25 hours per day and leisure falls by 1.04 hours, while market work increases slightly by 0.05 hours. Thus, under the strict categorization of time, the time input associated with young children is drawn from women's market work and both parents' leisure, with little change in men's market work. The time associated with childcare again declines as children grow older.[22]

Panel B reports per-child time inputs for the early 1990s, estimated as we described above. The 1990 estimates are smaller than their Panel A counterparts—for women, a child under five is associated with 1.76 hours of childcare per day in the early 1990s versus 2.18 in the ATUS years—consistent with the rise of intensive parenting documented by Ramey and Ramey (2010). They are also less precise, for two reasons. First, the pooled 1993–1995 surveys contain roughly 420 parent diaries for men and 620 for women, an order of magnitude fewer than the ATUS. Second, the imputation to overcome missing child ages compounds sampling error from both the AHTUS levels and the ATUS shape ratios. The imprecision is most severe for the work margin, where the underlying coefficients are small. Despite the differences in levels, the pattern across ages mirrors Panel A, where time inputs decline with child age, and men's contributions come almost entirely from leisure.

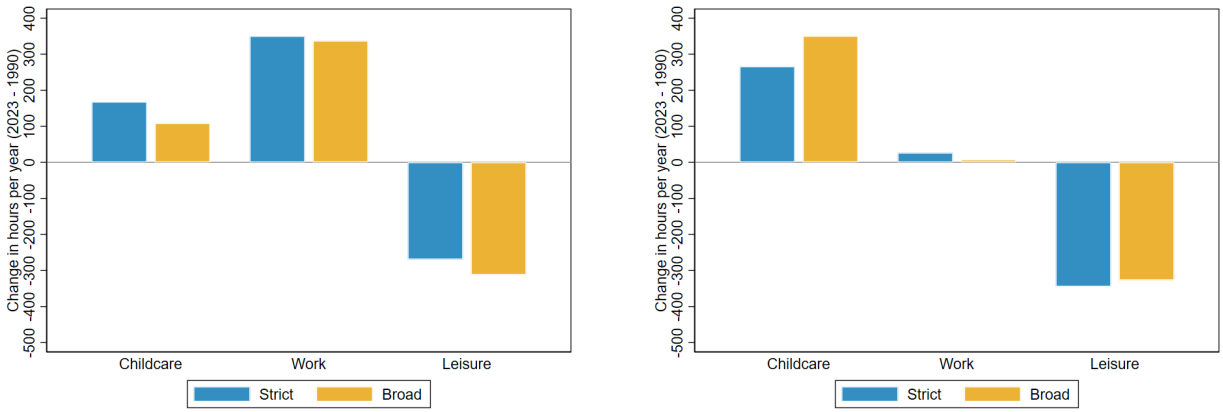
In Table 1, leisure pools sleep with actual leisure activity hours. Appendix Table B.1 separates sleep from leisure, revealing that, although parents sleep less than non-parents, the majority of the leisure decline is coming from a loss of actual leisure hours.

Appendix Table B.3 shows the same results using the secondary definition of time use. We again see that the correlation between children and the reallocation of time declines as children age, with the exception that men with teenagers work more. For robustness, Appendix Table B.4 shows the results from regressing activity hours on indicators for having the youngest child be 0-5, 6-11, and 12-17, as opposed to using counts. The patterns are similar.

Trends in childcare time. The pooled estimates in Table 1 describe average time reallocations over the pooled sample. For the CCI, however, we need the time inputs e_t^M and

[22]Housework is the excluded category from the table. For women, the measured declines in market work and leisure together exceed the measured rise in childcare time, meaning there is an increase in child-related home production (cooking, laundry, household tasks induced by children). Our baseline index prices the full work and leisure reductions, which implicitly treats this incremental home production as a child time cost. An alternative at the other end that rescales the priced hours to sum only to the childcare hours is discussed and shown to be nearly identical in Appendix E.3.

Figure 7: Changes in parental time use, 1990–2023



(a) Women

(b) Men

Notes: Bars show the change between 1990 and 2023 in the annual child penalty for each time-use category, measured in hours per child per year. The per-child time input in each year is the equal-weighted average across child age bins (0–5, 6–11, 12–17) of the coefficients from Equation 16, converted to annual hours. Positive bars indicate more hours devoted to the activity in 2023 than in 1990: the positive childcare bars show that childcare time per child has risen; the negative leisure bars show that parents give up more leisure than parents did in 1990; and the positive work bars show that parents give up fewer market-work hours. The *strict* measure counts only activities in which childcare is reported as the primary activity; the *broad* measure additionally includes time spent supervising or caring for children as a secondary activity while engaged in other tasks. Because the ATUS records secondary childcare only for children under 13, the broad measure for the 12–17 bin scales the strict estimate by the 6–11 broad-to-strict ratio. Estimates for 2003–2023 re-estimate Equation 16 year-by-year in the ATUS and are smoothed with a three-year moving average; the 1990 values are the imputed per-child inputs constructed from the pooled 1993–1995 time-diary surveys.

e_t^L to vary over time. We therefore re-estimate equation (16) allowing the coefficients on the number of children at each age to vary by year. Figure 7 plots the year-by-year child-induced time reallocations from 1990 to 2023, converted to annual hours.

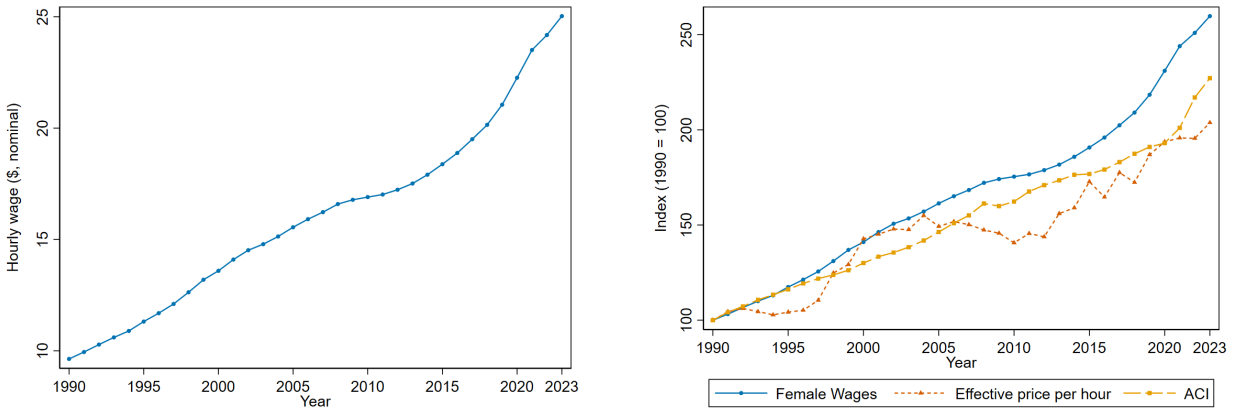
Consistent with the rise of intensive parenting documented by Ramey and Ramey (2010), childcare time among parents has increased substantially for both women and men. Under the strict definition, childcare hours rise by roughly 165 hours per year for women and 265 hours per year for men. But this increase does not come primarily from market work. For women, the market-work reduction associated with a child is *smaller* in 2023 than in 1990, while men experience essentially no change. However, the increase in childcare time for both men and women comes mainly from larger reductions in leisure. The leisure penalty grows by nearly 270 hours per year for women and 345 hours per year for men. Thus, the growing quantity margin in parental time is primarily a leisure margin, not a work margin.

Table 1: Impact of Children on Time Use (Primary)

	Women			Men		
	Childcare	Work	Leisure	Childcare	Work	Leisure
Panel A: ATUS pooled per-child effects, 2003–2023 (daily hours)						
Children 0-5	2.176***	-1.004***	-1.111***	1.247***	0.051	-1.040***
	(0.025)	(0.066)	(0.053)	(0.023)	(0.074)	(0.061)
Children 6-11	1.322***	-0.378***	-0.982***	0.822***	0.200***	-0.852***
	(0.024)	(0.065)	(0.053)	(0.022)	(0.077)	(0.063)
Children 12-17	0.848***	-0.139**	-0.868***	0.544***	0.297***	-0.737***
	(0.021)	(0.063)	(0.051)	(0.019)	(0.073)	(0.061)
Childless mean	0.15	4.41	15.65	0.07	5.43	15.73
Observations	75,498	75,498	75,498	65,123	65,123	65,123
Panel B: Imputed 1990 per-child inputs (daily hours)						
Children 0-5	1.763***	-1.747***	-0.696	0.625***	0.070	-0.478
	(0.173)	(0.410)	(0.431)	(0.123)	(0.696)	(0.724)
Children 6-11	0.829***	-1.084	-0.663	0.179	-0.018	-0.546
	(0.210)	(0.793)	(0.417)	(0.149)	(0.865)	(0.583)
Children 12-17	0.555***	-0.697	-0.610	0.108	0.644	-0.478
	(0.141)	(0.482)	(0.376)	(0.091)	(0.586)	(0.506)
Childless mean (AHTUS)	0.18	4.97	16.02	0.03	5.98	15.78
Observations (AHTUS)	2,281	2,281	2,281	1,974	1,974	1,974
Family-size	X	X	X	X	X	X
Age and education	X	X	X	X	X	X
Couple type	X	X	X	X	X	X
Day-of-week FE	X	X	X	X	X	X

Note: The sample consists of married and cohabiting couples and single adults aged 20–60, either childless (the omitted category) or with 1–3 own resident children aged 0–17. Each cell reports the coefficient on the number of the respondent's children in the indicated age bin from equation (16): a separate weighted OLS regression of daily hours in the indicated activity, by gender, with family-size indicators (two and three or more children), respondent age, education, and day-of-week fixed effects. Childcare uses the primary (strict) activity definition. Panel A is estimated on the pooled ATUS, 2003–2023, using ATUS final weights; robust standard errors in parentheses. Panel B reports imputed per-child inputs for the early 1990s: the coarse household-type coefficients estimated on the pooled 1993–1995 AHTUS surveys multiplied by the per-child-to-household shape ratios estimated in the early ATUS (2003–2007). Panel B standard errors are bootstrapped with 200 replications, resampling households with replacement within survey years, with the AHTUS and ATUS ingredients resampled independently and all quantities recomputed on each replicate. The childless mean is the weighted mean of daily hours among childless respondents in the relevant sample; AHTUS observations count person-diaries in the pooled 1993–1995 sample.

Figure 8: Female wages and the effective price of forgone work time



(a) Nominal hourly wage

(b) Wage, effective price, and ACI

Notes: Panel (a) shows the average nominal hourly wage for childless women aged 20–45, estimated from CPS/ACS data. We weight this sample to match the 1990 joint age×education distribution of women with children aged 0-6. Holding this composition fixed at 1990 ensures the series reflects wage growth within age–education cells rather than the changing composition of mothers. Panel (b) indexes the female wage, the effective price of forgone work time, and the Adult Consumption Index to 100 in 1990. The effective price per forgone work hour, w_t^{eff}, is defined as the total dollar earnings loss associated with motherhood divided by the total reduction in annual work hours.

Pricing work time. The time-use estimates show that the market-work margin is concentrated among women. We therefore price the work-time component of parental time using the earnings cost associated with motherhood.

A first benchmark is the market wage. Figure 8 documents the wage growth that underlies the opportunity cost of work time. Panel (a) shows the nominal hourly wage of married childless women aged 20–45. Panel (b) indexes this wage against the Adult Consumption Index constructed above. The nominal wages for women of child-bearing age rose substantially over the sample period, and wages also rose relative to the adult consumption benchmark.

The average wage is not, however, the right price for the work hours women give up when they become mothers. Motherhood does not simply remove hours from women's schedules at the prevailing wage. It may also shift women toward more flexible but lower-paying jobs, slow wage growth, reduce promotion probabilities, and, more broadly, generate earnings losses on inframarginal hours. The relevant price for the work-time margin is the total earnings cost per hour of market work forgone, not the average hourly wage of otherwise similar women.

We estimate this effective price by measuring the gap in work hours and earnings between married women with and without children. Using CPS/ACS data on married adults aged 20–45, we estimate how the presence and age of a first child affect annual earnings and

annual work hours. For outcome Y, we estimate

$$Y_{i,t}^{g} = \sum_{a=0}^{6} \beta_{a,t}^{g} \mathbf{1}\{\text{AgeOldestChild}_{i,t} = a\} + \gamma_{t}^{g} \text{NKids}_{i,t} + X_{i,t}' \theta_{t}^{g} + \varepsilon_{i,t}^{g}, \tag{17}$$

separately by gender g and separately for annual earnings and annual work hours. The omitted group is married individuals without children. The vector $X_{i,t}$ includes fully interacted race-by-education-by-state-by-parent age-by-metro status-by-calendar year fixed effects. We estimate the equation year-by-year, with a 5-year centered moving average applied to the level penalties and counterfactuals. Further details on constructing the penalty are in Appendix E.2.

For each year t, let $\widehat{\Delta E}_{t}^{f}$ denote the dollar earnings loss associated with motherhood and let $\widehat{\Delta H}_{t}^{f}$ denote the corresponding reduction in annual work hours. Both objects are constructed from equation (17), averaging the estimated child-age coefficients over ages zero through six. We then define the effective price of a forgone work hour as

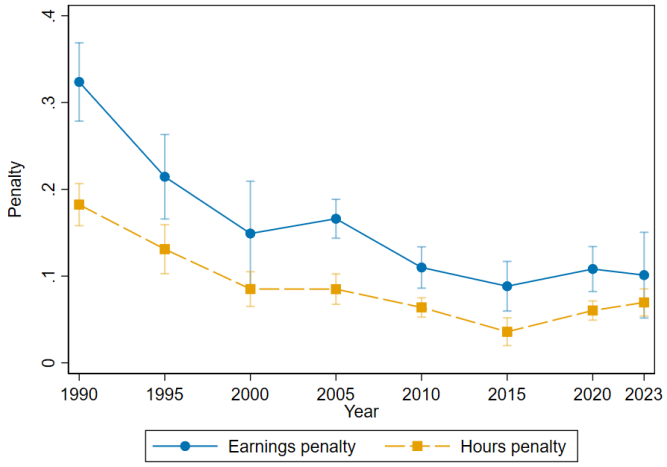
$$w_{t}^{\text{eff}} = \frac{\widehat{\Delta E}_{t}^{f}}{\widehat{\Delta H}_{t}^{f}}. \tag{18}$$

This price includes both the wage on hours not worked and any earnings losses on infra-marginal hours. Because annual hours include zeros for non-workers, $\widehat{\Delta H}_{t}^{f}$ captures both extensive-margin exits from employment and intensive-margin reductions in hours among employed women.

Figure 9 plots the estimated motherhood penalties over time. The earnings penalty for women falls substantially, from roughly 32 percent in 1990 to about 10 percent in 2023, consistent with broader convergence in women's labor-market outcomes. The hours penalty also falls over time, although less steeply than the earnings penalty. These estimates are consistent with the time-use evidence above: women still reduce market work when they have young children, but they do so by less than in earlier decades.[23] Appendix Figure A.9 separates the extensive- and intensive-margin labor-supply responses. Appendix Figure A.10 shows robustness to (1) including uninteracted fixed effects, (2) including all women/men regardless of marital status, (3) including children aged 1-17 and women/men up to age 60,

[23] Kleven (2025) reports a 10 percentage point decline in the child earnings penalty for women in the U.S. from 1990-2020. However, he measures the child penalty as the gap between men's and women's earnings change. Measuring it instead as the earnings gap between women with and without children produces a larger gap. The magnitudes are also similar across the CPS/ACS and time-use approaches. The estimates underlying Figure 7 imply that women worked 188 fewer hours per year when they had a young child in 1990 and 116 fewer hours in 2023. Assuming 50 work weeks per year, the CPS/ACS estimates in Figure 9 imply corresponding reductions of 214 and 129 hours.

Figure 9: Motherhood penalties in earnings and work hours



Notes: The figure plots the estimated percentage penalty in annual earnings and annual work hours for women, based on equation (17). Each series divides the estimated dollar (hours) penalty by counterfactual earnings (hours) of childless women, weighted to the 1990 joint age×education distribution of mothers. Level penalties and counterfactuals are smoothed with a five-year centered moving average. The figure shows the smoothed series at five-year intervals. Standard errors are clustered by state.

and (4) including 2-digit occupation fixed effects. For women, the levels of the penalty are different under these alternative specifications. For example, including 2-digit occupation fixed effects necessarily means we estimate on the sample of working women and the earnings and hours penalties are smaller. However, the decline in the penalties over time remains. For men, the results are stable with the exception of including all men, married or unmarried, in the sample. In this case, men with children have significantly higher earnings than men without children, pointing to the selection into marriage and having children documented in other research (Kleven et al., 2019).

This distinction has an important implication for the index. The average female wage rose substantially over the sample period, but the total earnings cost of motherhood grew more slowly because the motherhood earnings penalty declined and women now give up fewer work hours when they have children. Equivalently, w_t^{eff} generally exceeds the average hourly wage because it incorporates inframarginal career costs, but its growth is well below the growth of the average wage. Thus the work-time component of the CCI rises less than one would infer from wage growth alone.

Valuing leisure. The work-time calculation prices only childcare hours that displace market work. Table B.1 shows that this is an incomplete measure of the price of parental time, as a large share of the time parents devote to childcare comes from leisure rather than market

work. These hours do not generate an observed earnings loss, but they may nevertheless be costly because they displace time parents would otherwise have spent on non-market activities.

The difficulty is that leisure has no directly observed price. We therefore consider several alternatives. The most conservative case assigns zero value to displaced leisure time. This does not mean that leisure has no utility. Rather, it recognizes that measured reductions in leisure hours need not map one-for-one into forgone leisure services. Some childcare is joint with activities parents may still partially enjoy, such as eating, socializing, or watching television while supervising a child. In addition, leisure and work hours are not perfectly divisible: many parents cannot convert a marginal hour, or a set of fragmented minutes, into paid work. In such cases, valuing every displaced leisure hour at the wage may overstate the consumption-equivalent price. The zero-valuation case therefore provides a conservative lower bound that leaves forgone leisure unpriced.

A natural alternative benchmark comes from the standard labor-leisure model. If households can freely adjust work and leisure at the margin, the consumption-equivalent value of an additional hour of leisure equals the wage. This makes the source of the childcare hour irrelevant: a lost leisure hour and a lost work hour have the same marginal value. In the implementation below, we use w_t to value leisure hours, reserving w_t^{eff} for forgone work hours because w_t^{eff} also captures inframarginal earnings losses such as missed promotions, slower wage growth, and shifts toward more flexible jobs.

However, the wage-based benchmark may still miss changes over time in the value of leisure if leisure is itself produced with market goods. Many leisure activities combine time with complementary goods and services (e.g. an hour of downtime at home is better with TV and streaming services). A related literature emphasizes that improvements in leisure technologies and declines in the relative price of leisure goods can change the value of leisure time (Aguiar et al., 2021). Consistent with this, Appendix Figure A.7 shows that our price index for leisure goods and services has fallen by roughly 30 percent relative to the adult consumption index since 1990. If these goods and services are complementary to leisure time, then the consumption-equivalent value of a forgone leisure hour may rise even holding the wage fixed.

We summarize these cases with the following parameterization:

$$q_t^L = \phi \, w_t \left(\frac{\text{ACI}_t}{P_t^L} \right)^{\theta}, \tag{19}$$

where P_t^L is a price index for leisure goods and services and ACI_t is the adult-consumption index. The relative-price term is normalized to one in the base year. The parameter ϕ

governs the level of the leisure valuation relative to the wage, while θ governs how the value of leisure responds when leisure goods become cheaper or more expensive relative to adult consumption. Appendix D.4 provides a microfoundation for this functional form.

We report three main cases. The lower-bound case sets $\phi = 0$, so displaced leisure is not priced. The wage-valued case sets $\phi = 1$ and $\theta = 0$, so the value of leisure tracks the wage. The leisure-goods-adjusted case sets $\phi = 1$ and allows $\theta > 0$; at $\theta = 1$, a one-percent decline in the price of leisure goods relative to adult consumption raises the value of a leisure hour by one percent. We report results across these cases because the data identify the quantity of leisure displaced by children and the relative price of leisure goods, but they do not identify the shadow value of a forgone leisure hour.

With all three components priced, they fit naturally into a single extended Törnqvist index. Let G_t denote nominal spending on children's goods and define total child cost $\mathcal{C}_t \equiv G_t + w_t^{\text{eff}} e_t^M + q_t^L e_t^L$. The cost share of component $j \in \{G, M, L\}$ is $s_t^j = [\text{component } j\text{'s dollar value}]/\mathcal{C}_t$, with $s_t^G + s_t^M + s_t^L = 1$. The full Child Cost Index is then

$$\Delta \ln \text{CCI}_t = \bar{s}_t^G \, \Delta \ln \text{CCI}_t^g + \bar{s}_t^M \, \Delta \ln w_t^{\text{eff}} + \bar{s}_t^L \, \Delta \ln q_t^L, \tag{20}$$

where $\bar{s}_t^j = \frac{1}{2}(s_t^j + s_{t-1}^j)$ is the two-period average cost share. Equation (20) directly extends the goods-only formula (8): each term captures how much of total child costs a component represents and how fast its price is growing. When $\phi = 0$, the leisure share is zero and the index reduces to a goods-plus-work-time aggregator; it reduces further to CCI_t^g when forgone work hours are also negligible.

5.2 Results

Result 4: The forgone-work margin adds little to CCI growth. We first incorporate the market-work margin of parental time. Of course, adding mothers' forgone earnings raises the level of the price of a child in every year. The question for the price index is whether the work-time component became more expensive over time than the goods-and-services component or the adult benchmark. To answer this question, we construct a goods-plus-work index,

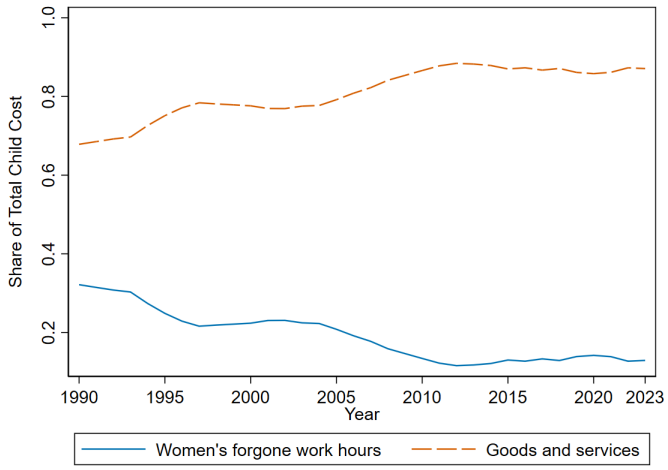
$$\Delta \ln \text{CCI}_t^{g+M} = \bar{s}_t^G \Delta \ln \text{CCI}_t^g + \bar{s}_t^M \Delta \ln w_t^{\text{eff}}, \tag{21}$$

where

$$s_t^M = \frac{w_t^{\text{eff}} e_t^M}{P_t^g G_t + w_t^{\text{eff}} e_t^M}$$

is the share of goods-plus-work child prices accounted for by mothers' forgone market work. The quantity of forgone work hours, e_t^M, affects the index through this cost share; the price

Figure 10: Price shares of goods spending and mothers' forgone work time



Notes: Each series shows the share of goods-plus-work child prices attributable to the indicated component. Total goods-plus-work price equals spending on children's goods and services plus the dollar value of mothers' forgone work earnings, $w_t^{\text{eff}} e_t^M$. The price index uses the corresponding Törnqvist two-period average shares.

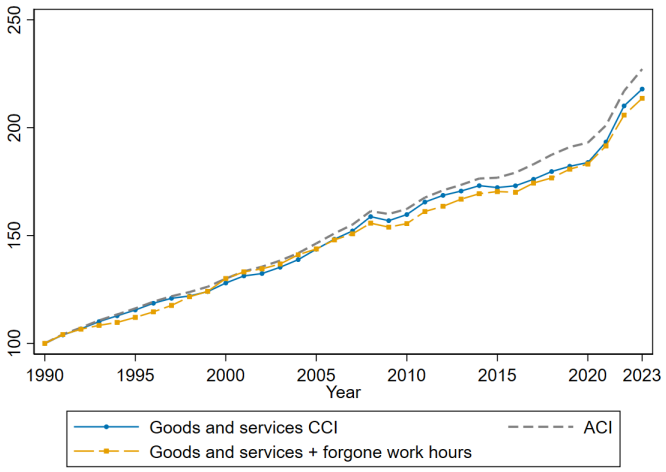
of those hours enters through w_t^{eff}.

Figure 10 shows the resulting price shares. In 1990, mothers' forgone work earnings accounted for roughly 33 percent of goods-plus-work child prices. By 2023, this share had fallen to about 13 percent. Two forces drive the decline. First, the motherhood earnings and hours penalties have both fallen substantially, so mothers now give up fewer market-work hours and a smaller fraction of earnings when they have children. Second, nominal spending on children's goods and services has grown, expanding the denominator of the price share. As a result, the goods-and-services share rises from roughly 65 percent to more than 80 percent over the sample period.

Figure 11 shows the resulting index. Adding mothers' forgone work earnings does not raise the growth of the CCI relative to the goods-only index. If anything, the goods-plus-work index lies slightly below the goods-only index for much of the sample and remains close to the adult benchmark. This does not mean that mothers' forgone earnings are unimportant in levels. Rather, it means that the work-time component has not become relatively more expensive than the goods-and-services bundle.

The intuition is visible in the components. The ordinary wage of childless women rose substantially over this period, but the effective price of a forgone work hour, w_t^{eff}, grew much more slowly than the ordinary wage because the motherhood earnings penalty declined. At the same time, the number of market-work hours displaced by children also fell. The first force limits the price growth of the work-time component; the second reduces its weight in

Figure 11: Goods-only CCI and CCI with mothers' forgone work time



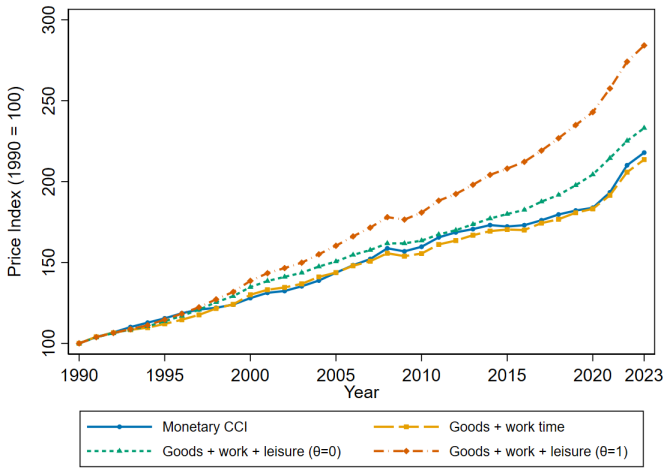
Notes: The goods-only CCI measures the price of the children's goods-and-services bundle, from equation (8). The goods-plus-work CCI adds mothers' forgone work earnings using equation (21). The adult benchmark is the Adult Consumption Index. All series are normalized to 100 in 1990. The figure excludes forgone leisure, which is incorporated in Result 5.

the index. Together, they imply that incorporating mothers' forgone market earnings leaves the relative price of children close to the goods-only benchmark.

Result 5: Forgone leisure is the consequential time margin. We finally add the leisure component of parental time. Because the shadow value of leisure is not observed, we report the range of valuations described in equation (19): leaving leisure unpriced, valuing it at the wage, and allowing its value to rise with the declining relative price of leisure goods. Figure 12 plots the resulting leisure-augmented indices. The goods-plus-work-time index, which leaves leisure unpriced, is the conservative lower-bound time-price measure from Result 4. When displaced leisure is valued at the wage, corresponding to $\phi = 1$ and $\theta = 0$, the index rises above both the goods-only CCI and the goods-plus-work-time CCI. The difference is economically meaningful because the quantity of leisure displaced by children is large. Even without any adjustment for leisure goods, pricing those hours adds a sizable time component that is absent from the work-time calculation.

The leisure-goods adjustment raises the index further. When $\theta = 1$, the value of a leisure hour adjusts one-for-one with the price of leisure goods relative to adult consumption. Since leisure goods and services have become cheaper relative to the adult consumption basket, this specification increases the value of each forgone leisure hour over time. The resulting index rises substantially faster than the wage-valued leisure index. Thus, the leisure margin is the component that can generate a meaningful rise in the relative price of a child. Across

Figure 12: Child Cost Index with forgone leisure



Notes: The figure plots four versions of the Child Cost Index, all normalized to 100 in 1990. The monetary CCI includes only children's goods and services. The goods-plus-work-time CCI adds mothers' forgone market-work earnings. The leisure-augmented indices additionally price childcare hours drawn from leisure using equation (19). The *valued at w* specification ($\phi = 1$, $\theta = 0$) values leisure at the wage. The *rising leisure value* specification ($\phi = 1$, $\theta = 1$) allows the value of leisure to adjust one-for-one with the price of leisure goods relative to adult consumption.

these specifications, the substantive conclusion is clear: forgone work earnings add little to relative child-price inflation, while forgone leisure is potentially much more important.

6 Conclusion

This paper constructs a Child Cost Index to measure whether raising children has become more expensive relative to adult consumption. The main takeaway is that the most visible child-related price increases are not the same as the price of the full child-rearing bundle. Childcare and tuition rose rapidly, but the goods-and-services CCI moved closely with the adult benchmark because those increases were offset by other parts of the child and adult baskets. Real spending per child also rose, but roughly in step with real spending on adults. Incorporating parental time changes the level and, depending on how leisure is valued, can raise the growth of the index, but the forgone-work margin alone adds little because mothers now give up fewer market-work hours than in the past. Taken together, the results suggest that the recent fertility decline is difficult to explain solely by a large increase in the average relative price of children.

References

Aguiar, M., M. Bils, K. K. Charles, and E. Hurst (2021): "Leisure luxuries and the labor supply of young men," *Journal of Political Economy*, 129, 337–382.

Bailey, M. J. (2006): "More power to the pill: The impact of contraceptive freedom on women's life cycle labor supply," *The Quarterly Journal of Economics*, 121, 289–320.

Becker, G. S. (1960): "An Economic Analysis of Fertility," 209–240.

——— (1965): "A Theory of the Allocation of Time," *The Economic Journal*, 75, 493–517.

Couillard, B. (2025): "Build, Baby, Build: How Housing Shapes Fertility," Tech. rep., Working Papers.

Deaton, A. and S. Zaidi (2002): "Guidelines for constructing consumption aggregates for welfare analysis," *The Review of Income and Wealth*, 48, 315–352.

Dow, A. (2026): "The Price of Parenthood: Childcare Costs and Fertility," Working paper, Boston University. Presented at SEHO (2025), ASHEcon (2025), APPAM (2025), and ASSA (2026).

Fernández, R., I. Berniell, and M. Onofri (2026): "Falling Fertility: The Changing Value of Freedom, Fulfillment, and Family," Tech. rep., National Bureau of Economic Research.

Galor, O. (2010): "Human capital, fertility and growth," in *Economic Growth*, Springer, 141–147.

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

Goldin, C. and L. F. Katz (2002): "The power of the pill: Oral contraceptives and women's career and marriage decisions," *Journal of Political Economy*, 110, 730–770.

Greenwood, J., A. Seshadri, and G. Vandenbroucke (2005): "The Baby Boom and Baby Bust," *American Economic Review*, 95, 183–207.

Hobijn, B., K. Mayer, C. Stennis, and G. Topa (2009): "Household Inflation Experiences in the U.S.: A Comprehensive Approach," Working Paper 2009-19, Federal Reserve Bank of San Francisco.

Jaravel, X. (2021): "Inflation Inequality: Measurement, Causes, and Policy Implications," *Annual Review of Economics*, 13, 599–629.

——— (2024): "Distributional consumer price indices," *Available at SSRN 5072082*.

Kearney, M. S., P. B. Levine, and L. Pardue (2022): "The puzzle of falling US birth rates since the great recession," *Journal of Economic Perspectives*, 36, 151–176.

Kleven, H. (2025): "The geography of child penalties and gender norms: A pseudo-event study approach," *American Economic Journal: Applied Economics*, Forthcoming.

Kleven, H., C. Landais, and J. E. Søgaard (2019): "Children and gender inequality: Evidence from Denmark," *American Economic Journal: Applied Economics*, 11, 181–209.

Laughlin, L. (2013): "Who's Minding the Kids? Child Care Arrangements: Spring 2011," Tech. Rep. P70-135, U.S. Census Bureau, Washington, DC.

Lino, M., K. Kuczynski, N. Rodriguez, and T. Schap (2017): "Expenditures on Children by Families, 2015," Miscellaneous publication no. 1528-2015, U.S. Department of Agriculture, Center for Nutrition Policy and Promotion, revised March 2017.

Miller, A. R. (2011): "The Effect of Motherhood Timing on Career Path," *Journal of Population Economics*, 24, 1071–1100.

Mincer, J. (1963): "Market Prices, Opportunity Costs, and Income Effects," 67–82.

New York Times/Morning Consult (2018): "Americans Are Having Fewer Babies. They Told Us Why," New York Times, survey of 1,858 adults conducted by Morning Consult for The New York Times.

OECD (2024): *Society at a Glance 2024: OECD Social Indicators*, Paris: OECD Publishing.

Pew Research Center (2024): "The Experiences of U.S. Adults Who Don't Have Children," Pew Research Center.

Ramey, G. and V. A. Ramey (2010): "The Rug Rat Race," *Brookings Papers on Economic Activity*, 41, 129–199.

Rothbarth, E. (1943): "The Measurement of Changes in Real Income under Conditions of Rationing," *The Review of Economic Studies*, 10, 100–107.

Santos, A., N. McGuckin, H. Y. Nakamoto, D. Gray, and S. Liss (2011): "Summary of Travel Trends: 2009 National Household Travel Survey," Technical report, U.S. Department of Transportation, Federal Highway Administration.

van Wijk, D. and P. Feijten (2025): "Rising House Prices, Falling Fertility? How Rising House Prices Widen Fertility Differences between Tenure Groups," *European Journal of Population*, 41, 33.

A Appendix Figures

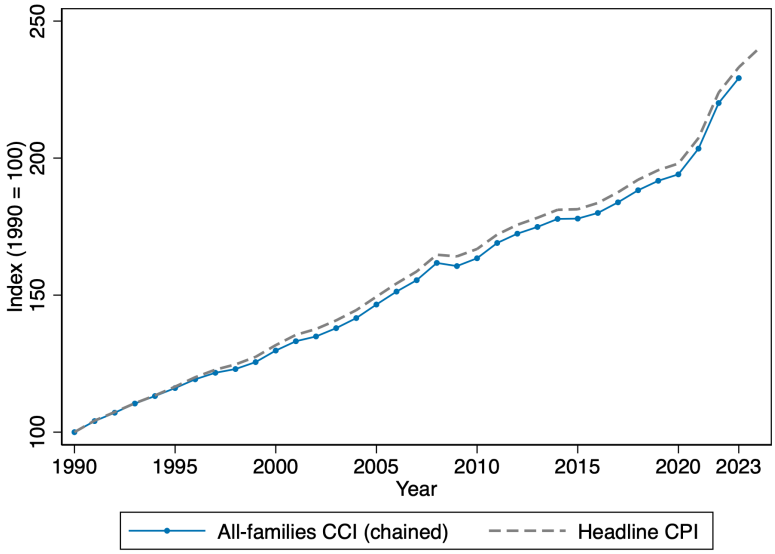


Figure A.1: Methodology validation: all-families price index versus BLS headline CPI

Notes: The figure plots a chained Törnqvist price index constructed using the expenditure weights of all families (including those with and without children) alongside the official BLS headline CPI, both normalized to 100 in 1990. Because the all-families weights span the full population, the resulting index should closely replicate the BLS headline CPI if the crosswalk from CEX expenditure classes to CPI subseries is correct. The close tracking between the two series confirms internal consistency of the methodology.

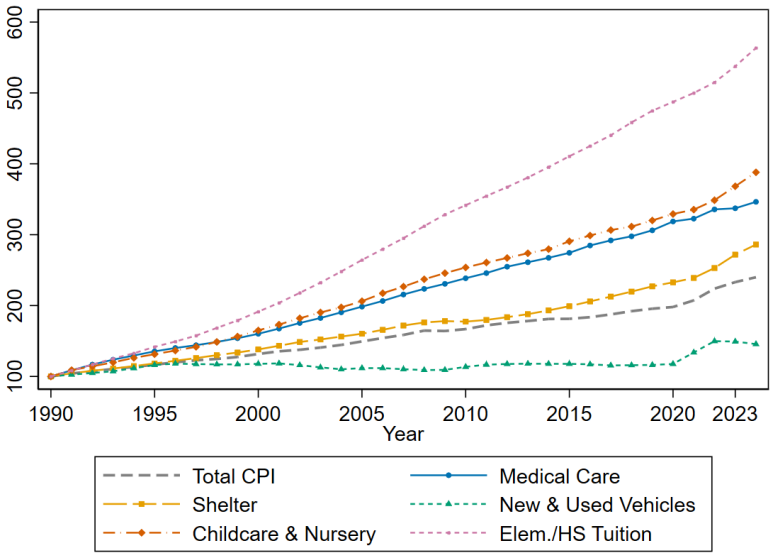
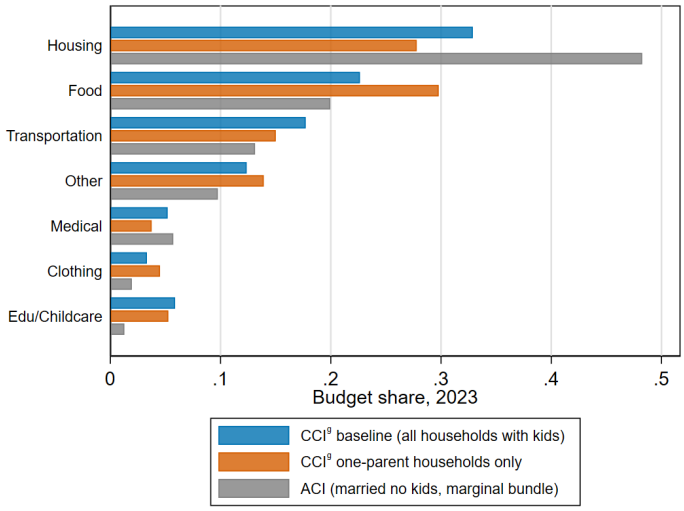


Figure A.2: Price growth of key spending categories relative to the headline CPI, 1990–2023

Notes: The figure plots price indices for five major spending categories alongside the BLS headline CPI (gray, dashed), all normalized to 100 in 1990. The categories shown are medical care (blue, solid), shelter (gold, long dash), new and used vehicles (green, short dash), childcare and nursery school—CPI code SEEB03 (orange, dash-dot)—and elementary and secondary school tuition—CPI code SEEB02 (pink, short dash). These categories are highlighted because they contribute the largest positive and negative components to the child–adult inflation gap documented in Figure 4: childcare and school tuition have risen far faster than overall prices, while vehicle prices have risen more slowly, and shelter has risen faster. The figure illustrates that the large positive and negative contributors to the decomposition each reflect genuine relative price movements in the underlying CPI sub-series, not artifacts of weighting. Source: Bureau of Labor Statistics, Consumer Price Index.

Figure A.3: Weights in single-parent and two-parent child basket vs. ACI



Notes: The figure compares expenditure weights in three 2023 consumption baskets: the child basket underlying the goods-and-services Child Cost Index (CCI^g, blue) of all households with kids, the CCI^g with only one parent households and the adult benchmark basket underlying the Adult Consumption Index (ACI, gray), where the ACI represents the marginal consumption bundle of married adults without children. The horizontal axis is each category's budget share. Shares are constructed from the CEX using the paper's allocation rules that map household expenditures to children.

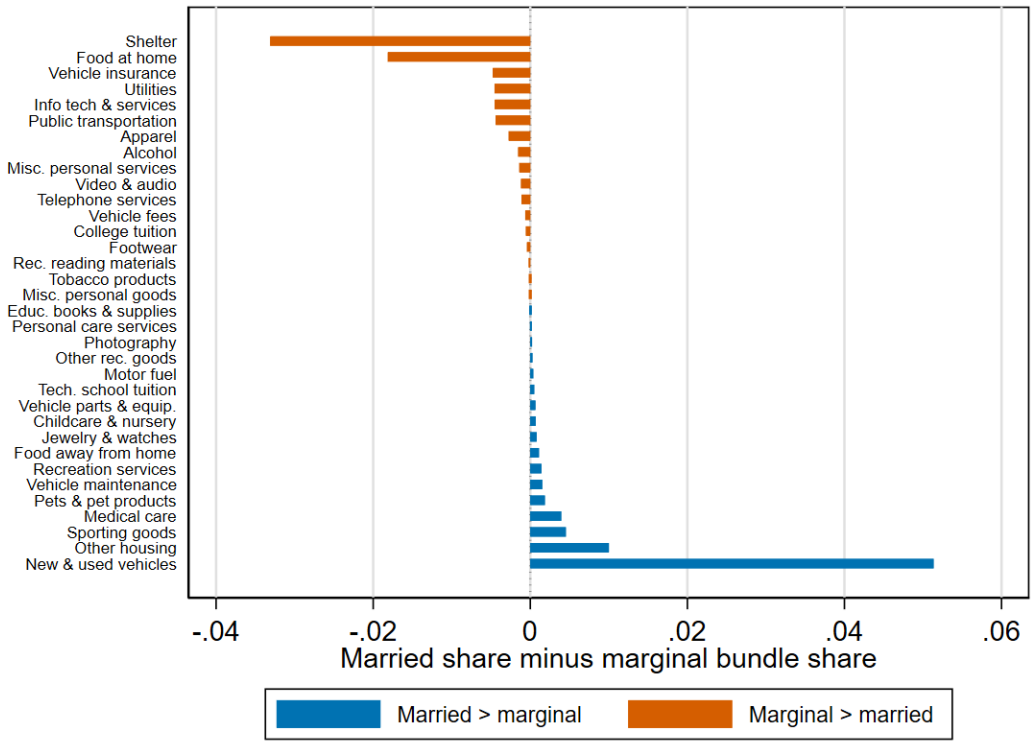


Figure A.4: Average versus marginal adult-consumption weights

Notes: The figure compares the average spending shares of married couples without children to the Engel-curve marginal weights used in the baseline Adult Consumption Index. The horizontal axis reports the average married-childless budget share minus the marginal ACI weight. Positive values indicate categories that receive more weight in the average married-childless bundle than in the marginal bundle; negative values indicate categories that receive more weight in the marginal bundle. Marginal weights are constructed from category-specific budget-share Engel curves estimated among married couples without children.

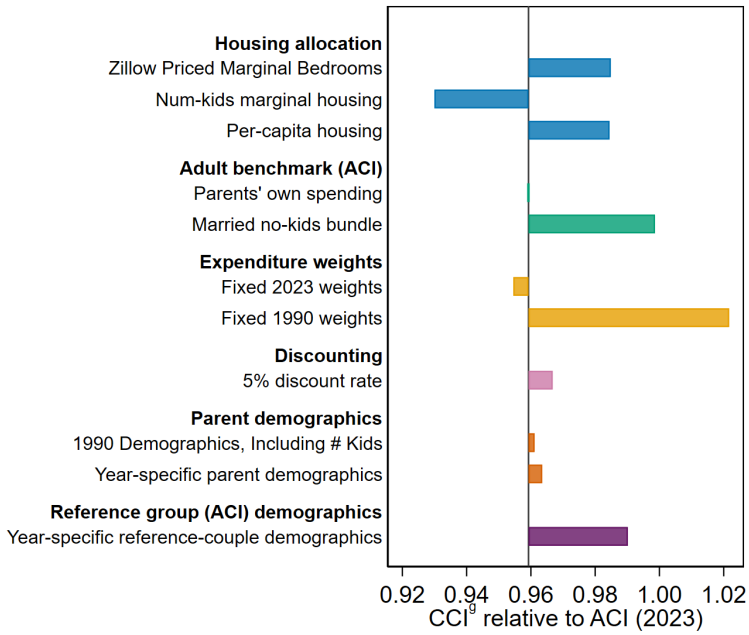


Figure A.5: Robustness of the 2023 relative price across specification choices

Notes: Each bar shows the 2023 ratio of the goods-and-services Child Cost Index (CCI^g) to the Adult Consumption Index (ACI) under an alternative specification. The solid vertical line at parity ($CCI^g/ACI = 1$) and the dashed line at the baseline ratio (bedroom housing, marginal ACI, chained weights) serve as reference points. Bars extending to the right of the dashed line indicate specifications that raise the estimated relative price of children above the baseline; bars to the left indicate lower relative prices. *Housing allocation*: per-capita division assigns all housing costs equally across household members; num-kids marginal uses observed differences in housing spending across family sizes. *Adult benchmark*: married no-kids bundle replaces the Engel-curve marginal bundle; parents' own spending uses consumption of parents net of child expenditures. *Expenditure weights*: fixed-weight alternatives use the 1990 or 2023 child basket in place of time-varying Törnqvist weights (each divided by the correspondingly weighted ACI). *Discounting*: applies a 5% annual discount rate to spending at later ages of childhood, upweighting early-childhood costs. *Parent demographics*: reweights the CEX sample to hold the joint distribution of parent age, education, race, and urban status fixed at 1990 levels. Source: Consumer Expenditure Survey and CPI.

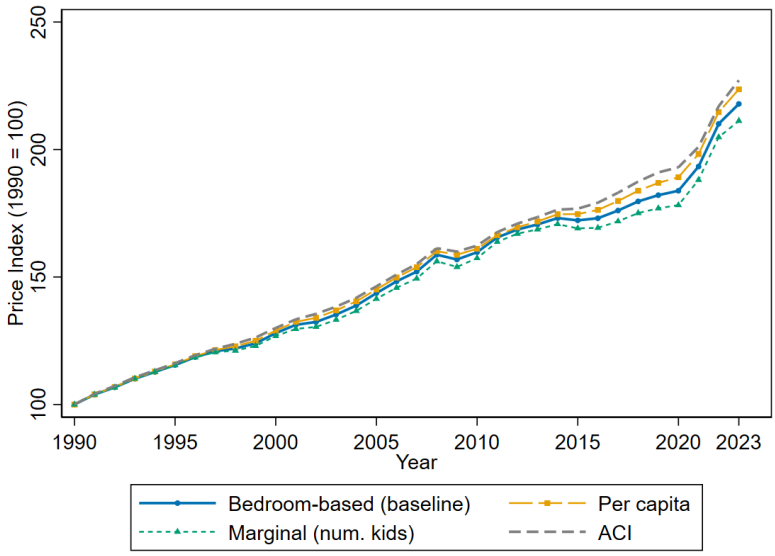


Figure A.6: Robustness to housing-allocation rule

Notes: The figure plots the chained goods-and-services Child Cost Index (CCI^g) under three approaches to allocating housing expenditures to children: the marginal-bedroom baseline (blue, solid), equal per-capita division (gold, long dash), and a comparison-based marginal cost estimated from differences in housing spending across family sizes (green, short dash). The Adult Consumption Index (ACI, gray dash) is shown for reference. All series are normalized to 100 in 1990. The close clustering of the three CCI series confirms that the main finding—the relative price of children's goods is stable—is not driven by the choice of housing-allocation rule.

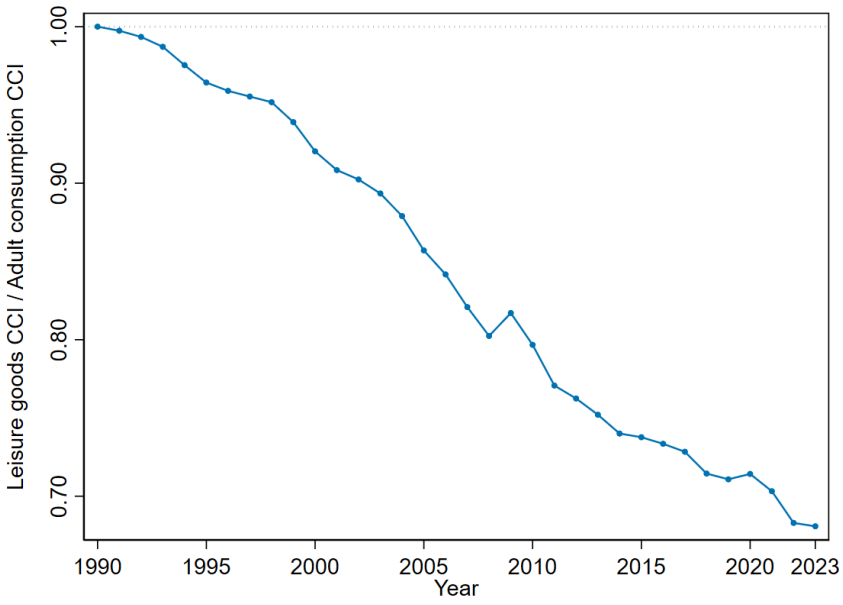


Figure A.7: The declining relative price of leisure goods

Notes: Ratio of a leisure-goods price index to the adult consumption index, both normalized to 100 in 1990. A ratio below one means leisure goods inflated more slowly than the overall adult basket.

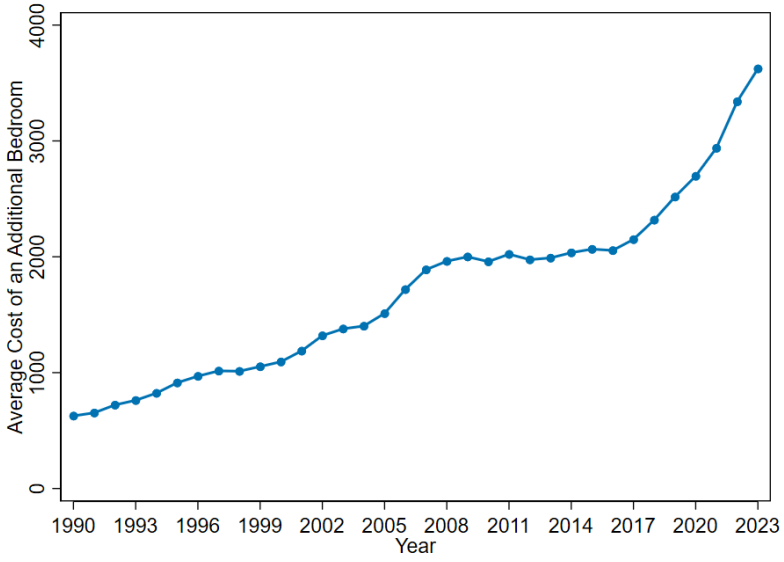


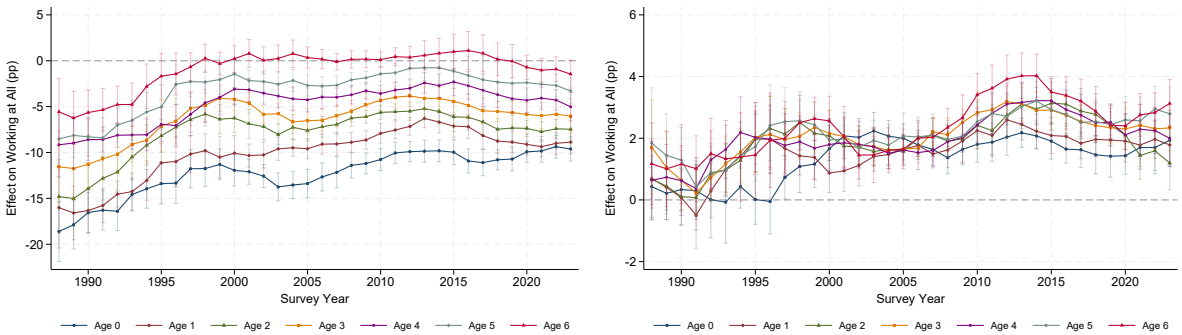
Figure A.8: Average Cost of an Additional Bedroom Over Time

Notes: This figure shows the average cost of an additional bedroom over time. For an additional bedroom k, the cost is $e^{\beta_0+\beta_k} - e^{\beta_0+\beta_{k-1}}$. The average cost is the unweighted average between the costs of a second bedroom, third bedroom, fourth bedroom, and fifth+ bedroom. This includes all families in the CEX data, not just families with children. Each year is in the given year's dollars.

Figure A.9: Intensive and Extensive Child Earnings and Hours Penalties

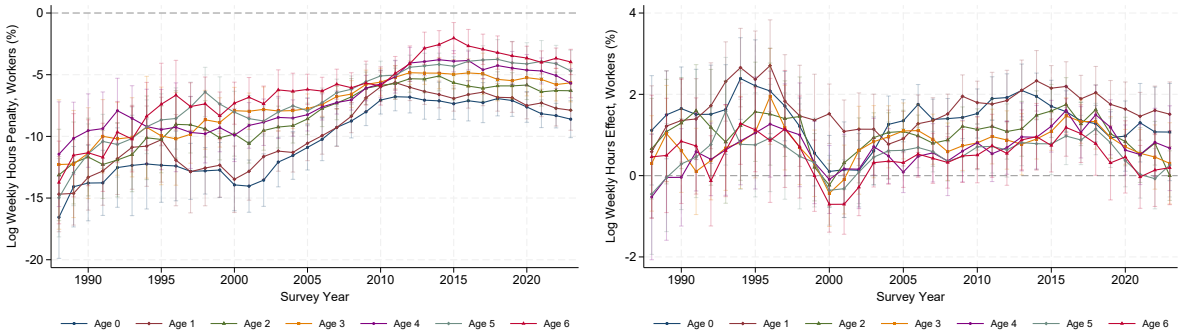
(a) **Extensive Margin Work: Women**

(b) **Extensive Margin Work: Men**



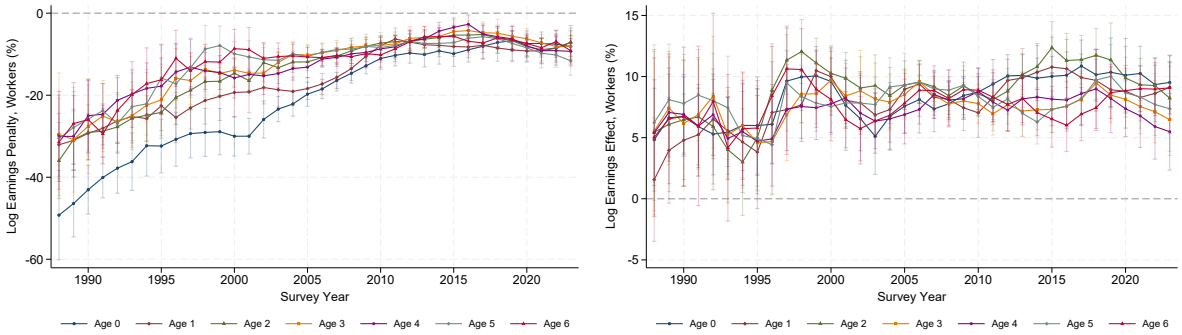
(c) **Intensive Margin Work: Women**

(d) **Intensive Margin Work: Men**



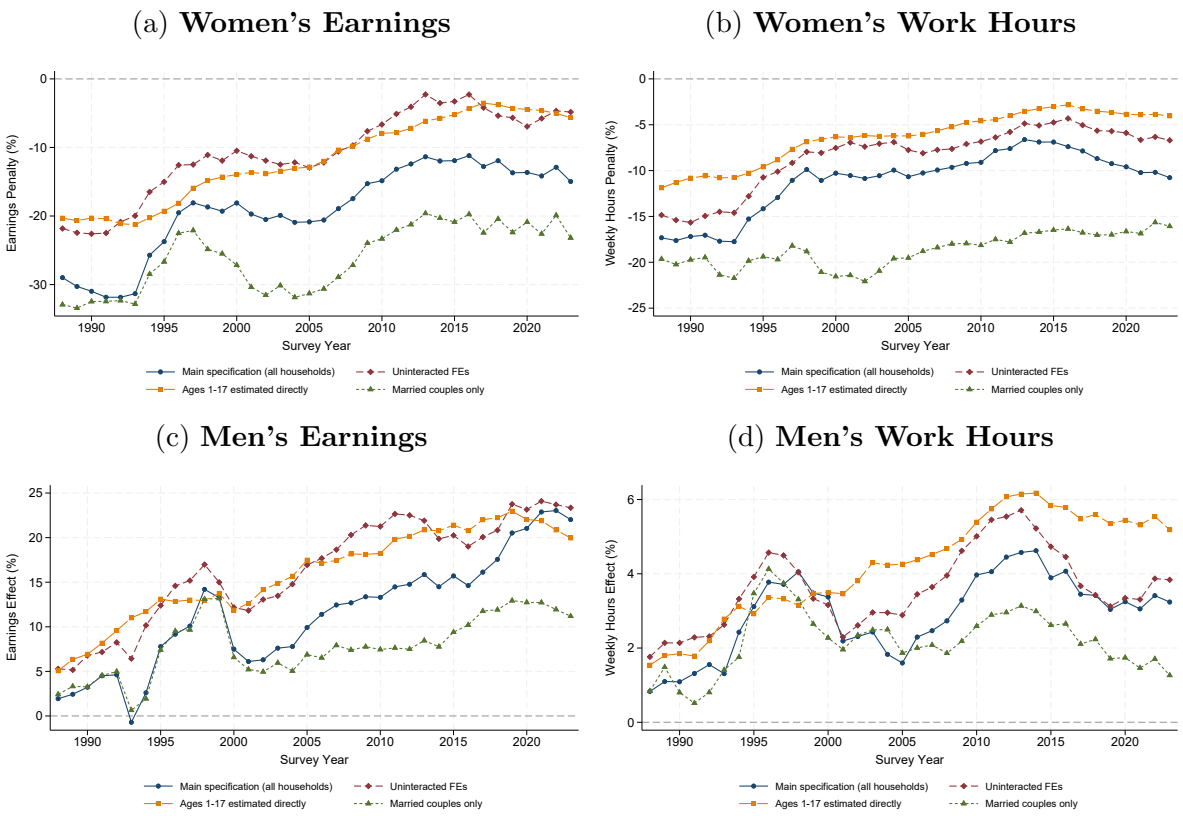
(e) **Intensive Margin Earnings: Women**

(f) **Intensive Margin Earnings: Men**



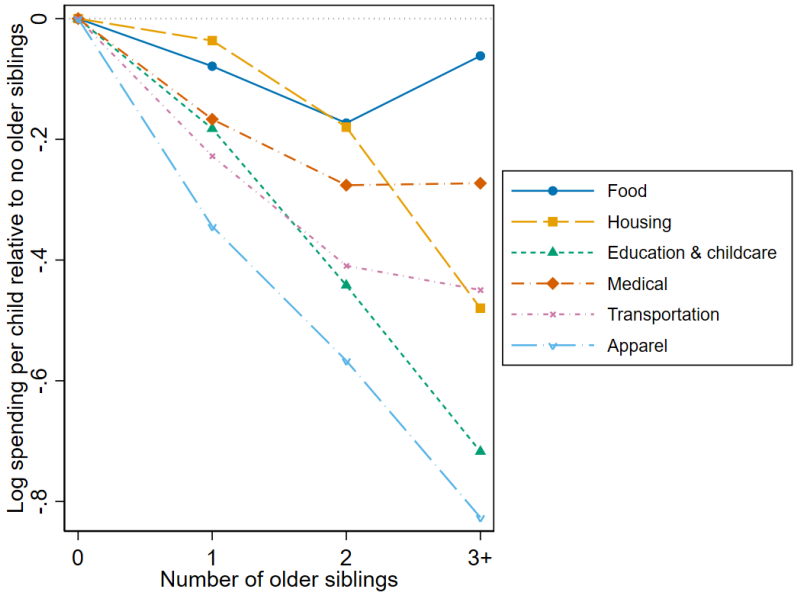
Note: Each panel plots child penalties by the age of the oldest child, on adults aged 20–45 in the pooled CPS ASEC and ACS, with childless individuals as the omitted group. Each series corresponds to a child age from zero to six; the horizontal axis is the income year. Panels (a) and (b) report the extensive margin: the percentage-point effect of a child of the indicated age on the probability of any market work, for women and men respectively. Panels (c) and (d) report the intensive margin in hours: the log weekly-hours penalty estimated on the sample of workers. Panels (e) and (f) report the intensive margin in earnings: the log earnings penalty, also estimated on workers. Because panels (c)–(f) condition on employment, they exclude the extensive-margin responses shown in panels (a) and (b). Estimates are computed in rolling five-year calendar-year windows. The vector of controls is fully interacted race-by-education-by-state-by-parent age-by-metro status-by-calendar year fixed effects; a survey fixed effect is included in years covered by both the CPS and the ACS, and weights are normalized to sum to one within survey×year. Vertical bars are 95 percent confidence intervals.

Figure A.10: Robustness of Child Earnings and Hours Penalties



Note: Each panel plots the child penalty in the indicated outcome under alternative specifications. Panels (a) and (b) report women's earnings and weekly work hours; panels (c) and (d) report the corresponding estimates for men. The horizontal axis is the income year. "Main specification" is the baseline used in our main analysis: adults aged 20–45 with an oldest child aged 0–6, fully interacted race-by-education-by-state-by-parent age-by-metro status-by-calendar year fixed effects. "Uninteracted FEs" replaces the interacted vector with additively separable fixed effects. "All women (incl. unmarried) / All men (incl. unmarried)" drops the marriage restriction. "Ages 1–17 estimated directly" extends the sample to parents whose oldest child is aged 1–17 and to adults up to age 60, estimating the penalty at every child age rather than extrapolating beyond age six. "2-digit occupation FEs (workers only)" adds two-digit occupation fixed effects and is therefore estimated on the sample of employed respondents, which removes the extensive margin and attenuates the level of the penalty. Series are smoothed with a three-year moving average.

Figure A.11: Spending by birth order: Category-level results



Notes: Each line plots log spending per child on a given category relative to spending on a child with no older siblings, as a function of the number of older siblings. Coefficients are estimated from a regression of log category spending per child on indicators for the number of older siblings, controlling for survey year, income decile, child age, urban status, number of adults in the household, and Census region. The sample is children aged 0–5 from the Consumer Expenditure Survey (CEX). Food and medicine show the smallest scale economies—indicating that these costs are nearly proportional to the number of children—while education and childcare and apparel fall sharply with sibling count, consistent with hand-me-downs and enrollment discounts.

B Appendix Tables

Table B.1: Impact of Children on Time Use (Separating sleep)

	Women				Men			
	Childcare	Work	Leisure	Sleep	Childcare	Work	Leisure	Sleep
Panel A: ATUS pooled per-child effects, 2003–2023 (daily hours)								
Children 0-5	2.176***	-1.004***	-1.024***	-0.087***	1.247***	0.051	-0.901***	-0.139***
	(0.025)	(0.066)	(0.046)	(0.033)	(0.023)	(0.074)	(0.053)	(0.034)
Children 6-11	1.322***	-0.378***	-0.883***	-0.100***	0.822***	0.200***	-0.778***	-0.074**
	(0.024)	(0.065)	(0.045)	(0.032)	(0.022)	(0.077)	(0.055)	(0.035)
Children 12-17	0.848***	-0.139**	-0.761***	-0.107***	0.544***	0.297***	-0.634***	-0.103***
	(0.021)	(0.063)	(0.043)	(0.032)	(0.019)	(0.073)	(0.053)	(0.033)
Childless mean	0.15	4.41	6.89	8.76	0.07	5.43	7.18	8.55
Observations	75,498	75,498	75,498	75,498	65,123	65,123	65,123	65,123
Panel B: Imputed 1990 per-child inputs (daily hours)								
Children 0-5	1.763***	-1.747***	-0.470	-0.079	0.625***	0.070	-0.267	–
	(0.173)	(0.410)	(0.320)	(0.115)	(0.123)	(0.696)	(0.816)	–
Children 6-11	0.829***	-1.084	-0.520	–	0.179	-0.018	-0.470	–
	(0.210)	(0.793)	(0.348)	–	(0.149)	(0.865)	(0.521)	–
Children 12-17	0.555***	-0.697	-0.483	–	0.108	0.644	-0.411	–
	(0.141)	(0.482)	(0.320)	–	(0.091)	(0.586)	(0.448)	–
Childless mean (AHTUS)	0.18	4.97	7.81	8.21	0.03	5.98	7.93	7.85
Observations (AHTUS)	2,281	2,281	2,281	2,281	1,974	1,974	1,974	1,974
Family-size controls	X	X	X	X	X	X	X	X
Age and education	X	X	X	X	X	X	X	X
Couple type	X	X	X	X	X	X	X	X
Day-of-week FE	X	X	X	X	X	X	X	X

Note: The sample consists of married and cohabiting couples and single adults aged 20–60, either childless (the omitted category) or with 1–3 own resident children aged 0–17. This version replicates Table 1 but separates sleep from the leisure category. Each cell reports the coefficient on the number of the respondent's children in the indicated age bin from equation (16): a separate weighted OLS regression of daily hours in the indicated activity, by gender, with family-size indicators (two and three or more children), respondent age, education, and day-of-week fixed effects. Childcare uses the primary (strict) activity definition. Panel A is estimated on the pooled ATUS, 2003–2023, using ATUS final weights; robust standard errors in parentheses. Panel B reports imputed per-child inputs for the early 1990s: the coarse household-type coefficients estimated on the pooled 1993–1995 AHTUS surveys multiplied by the per-child-to-household shape ratios estimated in the early ATUS (2003–2007). Panel B standard errors are bootstrapped with 200 replications, resampling households with replacement within survey years, with the AHTUS and ATUS ingredients resampled independently and all quantities recomputed on each replicate. Missing coefficients were not imputed because the associated early ATUS(2003-2007) effect was close to 0. The childless mean is the weighted mean of daily hours among childless respondents in the relevant sample; AHTUS observations count person-diaries in the pooled 1993–1995 sample.

Table B.2: Impact of Children on Time Use, AHTUS

	Women			Men		
	Childcare	Work	Leisure	Childcare	Work	Leisure
Youngest under 5	2.092***	-2.120***	-0.804*	0.716***	0.199	-0.563
	(0.205)	(0.493)	(0.458)	(0.139)	(0.718)	(0.662)
Youngest 5–17	0.574***	-0.684	-0.645	0.128	0.714	-0.564
	(0.144)	(0.461)	(0.412)	(0.099)	(0.641)	(0.573)
Observations	2,281	2,281	2,281	1,974	1,974	1,974

Note: The sample is the AHTUS (1993–1995; 1992 dropped), showing a regression of daily hours on the two household-type indicators (youngest child under 5, youngest 5–17) plus family-size controls, age, education, and day-of-week FE; adults 20–60 in childless households or households with 1–3 children. Robust standard errors in parentheses.

Table B.3: Impact of Children on Time Use (Secondary)

	Women			Men		
	Childcare	Work	Leisure	Childcare	Work	Leisure
Panel A: ATUS pooled per-child effects, 2003–2023 (daily hours)						
Children 0-5	7.011***	-1.239***	-3.785***	4.537***	-0.117	-3.199***
	(0.053)	(0.065)	(0.051)	(0.054)	(0.075)	(0.058)
Children 6-11	6.087***	-0.676***	-3.701***	4.476***	-0.006	-3.255***
	(0.054)	(0.065)	(0.049)	(0.058)	(0.077)	(0.060)
Children 12-17†	3.369***	-0.294***	-2.363***	2.436***	0.191***	-2.020***
	(0.049)	(0.062)	(0.051)	(0.048)	(0.073)	(0.060)
Childless mean	0.21	4.41	15.61	0.15	5.43	15.68
Observations	75,498	75,498	75,498	65,123	65,123	65,123
Panel B: Imputed 1990 per-child inputs (daily hours)						
Children 0-5	6.644***	-1.934***	-3.384***	3.853***	-0.058	-2.663***
	(0.184)	(0.413)	(0.440)	(0.147)	(0.695)	(0.731)
Children 6-11	5.590***	-1.311*	-3.419***	3.797***	-0.183	-2.965***
	(0.226)	(0.796)	(0.418)	(0.172)	(0.863)	(0.583)
Children 12-17†	3.739***	-0.843*	-3.144***	2.295***	–	-2.594***
	(0.194)	(0.487)	(0.450)	(0.157)	–	(0.606)
Observations (AHTUS)	2,281	2,281	2,281	1,974	1,974	1,974
Family-size controls	X	X	X	X	X	X
Age and education	X	X	X	X	X	X
Couple type	X	X	X	X	X	X
Day-of-week FE	X	X	X	X	X	X

Note: This table replicates Table B.1 but uses the expansive definition of childcare. Because the AHTUS does not measure secondary time use, each Panel B coefficient is constructed as the full-history primary coefficient from Panel A plus the expansive-minus-primary gap estimated on ATUS, projected back over the AHTUS years. Panel B standard errors (in parentheses) are bootstrapped with 200 replications, resampling households with replacement within survey years, with all ingredients computed on the same resample. Missing coefficients were not imputed because the associated early ATUS(2003-2007) effect was close to 0. Secondary childcare is recorded only for children under 13, so the youngest 12–17 expansive gap cannot be estimated directly; it is imputed by scaling the 12–17 primary coefficient by the 6–11 expansive/primary ratio.

Table B.4: Impact of Children on Time Use

	Women			Men		
	Childcare	Work	Leisure	Childcare	Work	Leisure
Panel A: Primary Definition (estimated, AHTUS+ATUS)						
Youngest child 0-5	2.486***	-1.499***	-1.198***	1.384***	0.296***	-1.439***
	(0.032)	(0.085)	(0.073)	(0.028)	(0.089)	(0.078)
Youngest child 6-11	1.154***	-0.504***	-0.983***	0.712***	0.529***	-1.128***
	(0.030)	(0.089)	(0.073)	(0.027)	(0.092)	(0.080)
Youngest child 12-17	0.492***	-0.099	-0.851***	0.285***	0.584***	-0.874***
	(0.023)	(0.085)	(0.071)	(0.017)	(0.099)	(0.081)
Control Mean	0.15	4.52	15.74	0.07	5.49	15.77
Observations	77751	77751	77751	67115	67115	67115
Panel B: Expansive Definition (primary + ATUS gap projected over AHTUS years)						
Youngest child 0-5	8.693***	-1.793***	-4.599***	5.579***	0.086	-4.203***
	(0.065)	(0.081)	(0.066)	(0.067)	(0.099)	(0.079)
Youngest child 6-11	6.709***	-0.869***	-4.156***	4.933***	0.292***	-3.935***
	(0.063)	(0.092)	(0.067)	(0.065)	(0.097)	(0.080)
Youngest child 12-17†	2.859***	-0.112	-3.430***	1.953***	0.316***	-2.916***
	(0.113)	(0.153)	(0.240)	(0.117)	(0.102)	(0.233)
Control Mean (ATUS)	0.21	4.41	15.61	0.15	5.43	15.68
Observations (gap)	75498	75498	75498	65123	65123	65123

Note: The sample consists of married and cohabiting couples and single adults aged 20–60, either childless (the omitted category) or with 1–3 own resident children aged 0–17; AHTUS 1993–1995 pooled with ATUS 2003–2023, weighted by tufnwgtp (AHTUS weights rescaled to ATUS year totals). Each column reports coefficients on indicators for the age group of the household's youngest child from a regression of daily hours on the three household-type indicators, family-size controls (two- and three-child indicators), respondent age, education, and day-of-week fixed effects. Panel A uses the primary activity definition and is estimated directly on the pooled AHTUS+ATUS sample; robust standard errors in parentheses. Panel B reports expansive-definition effects (secondary childcare reclassified from leisure, work, and home production into childcare); because the AHTUS secondary-childcare measure is not comparable to the ATUS measure, each Panel B coefficient is constructed as the full-history primary coefficient from Panel A plus the expansive-minus-primary gap estimated on ATUS (2003–2023), projected back over the AHTUS years. Panel B standard errors (in parentheses) are bootstrapped with 200 replications, resampling households with replacement within survey years, with all ingredients computed on the same resample. Secondary childcare is recorded only for children under 13, so the youngest 12–17 expansive gap cannot be estimated directly; it is imputed by scaling the 12–17 primary coefficient by the 6–11 expansive/primary ratio.

C Family-Specific Housing Price Growth

In this appendix, we investigate two possible sources of such differential inflation in the housing that couples with and without children consume: (1) the differential price growth trends for housing in high-quality public school districts, and (2) the differential price trends by the size of the housing unit.

The distinction between price *levels* and price *growth* is central to these exercises. Homes in high-performing school districts and homes with more bedrooms are more expensive in every year. Those level premiums may affect the amount of resources required to raise children. They affect the growth of a fixed-quality Child Cost Index, however, only if the relative price of those housing attributes increases over time. Throughout, we will find little evidence that the relative price of family-specific housing has increased in recent years.

C.1 School quality and housing costs

Access to public schools is bundled with residential location. Broad shelter indices could therefore understate family-specific inflation if the cost of living in a high-quality school district rose systematically faster than the cost of living in other districts.

We examine this possibility by combining school-district housing outcomes from the American Community Survey (ACS) with district-level test scores from the Stanford Education Data Archive (SEDA). The housing outcomes are median owner-occupied home value and median gross rent. Because the ACS publishes these district-level outcomes as five-year estimates, we assign each observation to the final year of its five-year window. We pool annual SEDA scores over the corresponding trailing five years and merge the two sources using districts' NCES identifiers. The owner-occupied sample covers 2013–2023, while the rental sample covers 2011–2023.

In our baseline specification, we rank districts by test scores within state in each year and divide them into the top 10 percent, the middle 10 percent (the 45th–55th percentiles), and the bottom 10 percent. The composition of each group is therefore allowed to change over time. This construction answers how the price of housing in a district that is currently near the top of its state's school-quality distribution has evolved relative to the price of housing in a middle- or bottom-ranked district. We find similar results when we fix the ranking of districts in the initial year.

Figure C.1 reports the results. Panels (a) and (c) show that school quality is strongly associated with housing-cost *levels*. In 2013, for example, the average district-level median home value in the top test-score decile was nearly three times that in the bottom decile. The corresponding rent gap is smaller but also substantial. Unsurprisingly, higher-performing

public schools carry a large housing premium at a point in time.

The growth comparisons, however, do not show that this premium has systematically widened. Panel (b) shows that owner-occupied home values rose by 52 percent in the top-decile group between 2013 and 2023. This was less than the 64 percent increase in the middle group and the roughly 59 percent increase in the bottom group. All three house price index series rose faster than shelter CPI over this window.

Panel (d) provides the breakdown by school district using rents, which is a more direct comparison to shelter within the CPI. From 2011 to 2023, median gross rent rose by 47 percent in the top-decile group, 45 percent in the middle group, and 44 percent in the bottom group. The top-to-bottom difference is only 3.1 percentage points over twelve years, and rent growth in all three groups was just below the 51 percent increase in shelter CPI. Overall, these patterns show no indication that applying the aggregate shelter index causes us to understate the price growth faced by families seeking access to currently high-performing districts.

C.2 Housing size and bedroom premiums

A different reason that families may be facing a higher price index is that families disproportionately consume larger homes. Even if families require only the incremental space associated with an additional bedroom, broad shelter indices could understate this cost if three- and four-bedroom units have appreciated faster than smaller units.

We investigate this possibility using the Zillow Home Value Index by bedroom count for cities and ZIP codes from 2000 through 2023. For an outcome $X \in \{V, R\}$—home value V or rent R—define the premium associated with moving from a $(b-1)$-bedroom unit to a b-bedroom unit in area a and year t as

$$p^X_{a,b,t} = \frac{X_{a,b,t}}{X_{a,b-1,t}} - 1. \tag{C.1}$$

For example, $p^V_{a,3,t} = 0.40$ means that a three-bedroom home is valued 40 percent above a two-bedroom home in the same area. A decline in this premium between two dates means that the larger unit appreciated less rapidly than the adjacent smaller unit. We compute the premium within each geography and report population-weighted averages across those locations.[24] The comparisons within city allow for large houses to be located in different

[24]Because bedroom-specific coverage varies substantially, particularly for one-bedroom units, we use a pairwise sample for each comparison (e.g. the 1-to-2 bedroom comparison requires only the one- and two-bedroom series and the 2-to-3 comparison requires only the two- and three-bedroom series). At the city level, the resulting samples contain 678, 3,177, and 3,449 cities. At the ZIP-code level, they contain 1,429, 7,999, and 8,583 ZIP codes, respectively.

neighborhoods than smaller units (e.g. it allows comparisons between urban areas and suburban areas), while the zip-code level comparisons look only at price differentials between different sized units in the same neighborhood.

Panels (a) and (c) of Figure C.2 compare the home-value premiums in 2000 and 2023 within city and within a zip code, respectively. Unsurprisingly, all of the premiums are positive on average, meaning that a marginal bedroom increases the home's value. However, more importantly, the 2023 bar is lower than the 2000 bar in all cases. Across cities, the second-bedroom premium falls from about 40 percent to 34 percent, the third-bedroom premium declines from about 35 percent to 33 percent, and the fourth-bedroom premium falls from about 35 percent to 30 percent. Across ZIP codes, the corresponding premiums fall from about 54 to 46 percent, from 42 to 39 percent, and from 31 to 28 percent. Thus, in the Zillow home price data, family-sized homes did not appreciate faster than smaller homes in the same area. If anything, their relative price declined.

As we discussed above, home values are asset prices and need not move one-for-one with the flow price of shelter services. We therefore repeat the exercise using ACS median gross rent by bedroom count. These data are available annually at the state and county levels beginning in 2015.[25] Panels (b) and (d) show that the results using rents reinforce the home-value evidence. At the state level, the second-bedroom premium falls from about 21 percent in 2015 to 16 percent in 2023, and the third-bedroom premium falls from about 20 percent to 14 percent.

We acknowledge that the bedroom-specific series do not cover the full 1990–2023 CCI window, and neither the ACS medians nor the Zillow indices provide a hedonic price for an otherwise identical unit with one additional bedroom. Nevertheless, the direction of the evidence is at least suggestive that these patterns would not meaningfully affect our conclusions about the movements in the CCI relative to the ACI.

D Conceptual Framework Details

This appendix provides the algebra underlying the conceptual framework in Section 2. The appendix has three goals. First, it derives the fixed-quality cost function that motivates the Child Cost Index. Second, it shows how the share-weighted index used in the data follows from that cost function. Third, it explains why the relevant price for fertility decisions is the Child Cost Index relative to the adult consumption index.

[25]The county pairwise samples contain 800, 799, and 742 counties for the second-, third-, and fourth-bedroom comparisons, respectively.

D.1 Household problem

For notational simplicity, we suppress time subscripts in this section. The household has preferences over adult consumption c and child services Z:

$$U(c, Z) = \left[(1-\omega)c^{\rho} + \omega Z^{\rho}\right]^{1/\rho}, \qquad \rho = \frac{\sigma - 1}{\sigma}, \tag{D.2}$$

where σ is the elasticity of substitution between adult consumption and child services, and $\omega \in (0, 1)$ governs the relative weight on children. Child services combine the number of children n and the quality, or resource intensity, of each child q:

$$Z = nq^{\eta}, \qquad \eta \in (0, 1). \tag{D.3}$$

We treat n as continuous for the purpose of deriving the household's demand conditions.

Quality is produced using parental time e and market goods x:

$$q = Ae^{\alpha}x^{1-\alpha}, \qquad \alpha \in (0, 1), \tag{D.4}$$

where A is fixed production technology. Each child also requires fixed care time h, independent of quality. This fixed component is not needed to derive the fixed-quality price index, but it gives the quantity–quality problem an interior solution by making each additional child require some resources even before quality investments are chosen.

The household has time endowment $\overline{T}$, non-labor income V, wage w, price of adult consumption P^a, and price of child market goods p. Its full-income budget constraint is

$$P^a c + n\left[px + w(h + e)\right] = w\overline{T} + V \equiv I. \tag{D.5}$$

The term in brackets is the resource cost of one child. It includes market spending on child goods, px, fixed care time, wh, and quality-producing parental time, we.

We solve the household problem in two stages. First, conditional on producing a child of quality q, the household chooses the cost-minimizing combination of time and goods. Second, given the resulting cost function, the household chooses adult consumption, the number of children, and child quality. The first stage delivers the Child Cost Index; the second stage, developed below, clarifies why fertility depends on the relative price of children.

D.2 Cost minimization and the fixed-quality cost of a child

Conditional on producing quality q, the household chooses the cost-minimizing input mix:

$$\min_{e,x} \left\{ w(h+e) + px : Ae^{\alpha}x^{1-\alpha} \geq q \right\}. \tag{D.6}$$

The fixed care cost wh does not affect the choice of e and x, so the first-order conditions for the variable inputs are

$$w = \lambda \alpha A e^{\alpha-1} x^{1-\alpha}, \qquad p = \lambda(1-\alpha) A e^{\alpha} x^{-\alpha}.$$

Dividing the two conditions gives the cost-minimizing input ratio:

$$\frac{e^*(q;w,p)}{x^*(q;w,p)} = \frac{\alpha}{1-\alpha}\frac{p}{w}. \tag{D.7}$$

When parental time becomes more expensive relative to child goods, the household substitutes away from time and toward goods in producing a given level of quality.

Substituting this input ratio into the quality constraint gives the dual unit cost of child quality:

$$\pi^q(w,p;A) = B(A) w^{\alpha} p^{1-\alpha}, \qquad B(A) = A^{-1}\alpha^{-\alpha}(1-\alpha)^{-(1-\alpha)}. \tag{D.8}$$

The term $\pi^q(w,p;A)$ is the minimum variable cost of producing one unit of child quality. Since the Cobb–Douglas technology is constant returns to scale in e and x, the variable cost of producing quality q is $\pi^q(w,p;A)q$. The full cost of one child of quality q is therefore

$$\mathcal{C}(q;w,p) = wh + \pi^q(w,p;A)q. \tag{D.9}$$

Equivalently, evaluated at the cost-minimizing inputs,

$$\mathcal{C}(q;w,p) = w\left[h + e^*(q;w,p)\right] + px^*(q;w,p). \tag{D.10}$$

The Child Cost Index is the price index associated with this cost function evaluated at a fixed reference quality $\bar{q}$:

$$\mathrm{CCI}(w,p) \equiv \mathcal{C}(\bar{q};w,p). \tag{D.11}$$

The reference quality $\bar{q}$ is held fixed by construction. The cost-minimizing input mix may change as relative prices change, but the object being priced is always the same child quality.

Holding $\bar{q}$ and A fixed, Shephard's lemma implies that the elasticity of the fixed-quality

cost function with respect to each input price equals the corresponding cost share:

$$\frac{\partial \ln \mathrm{CCI}}{\partial \ln p} = s_G \equiv \frac{p x^*(\bar{q}; w, p)}{\mathrm{CCI}}, \qquad \frac{\partial \ln \mathrm{CCI}}{\partial \ln w} = s_T \equiv \frac{w\,[h + e^*(\bar{q}; w, p)]}{\mathrm{CCI}}, \tag{D.12}$$

with $s_G + s_T = 1$. The fixed care component h enters the time share s_T; when $h = 0$, the time and goods shares reduce to the Cobb–Douglas elasticities α and $1 - \alpha$.

The growth rate of the fixed-quality child cost is therefore

$$d \ln \mathrm{CCI} = s_G d \ln p + s_T d \ln w. \tag{D.13}$$

This expression is the theoretical counterpart of the share-weighted index constructed in the data. We implement it using a Törnqvist index, which weights input-price growth by average expenditure shares across adjacent periods.

D.3 The relative price of children and fertility demand

The previous subsection derives the fixed-quality cost of a child. We next show why the relevant price for the household's fertility decision is this child cost relative to the price of adult consumption.

For a child of reference quality $\bar{q}$, one child produces $\bar{q}^\eta$ units of child services. The price of one unit of fixed-quality child services is therefore

$$\pi^Z(\bar{q}; w, p) = \frac{\mathrm{CCI}(w, p)}{\bar{q}^\eta}. \tag{D.14}$$

Since $\bar{q}$ is fixed by construction, changes in π^Z are exactly changes in the CCI. The relative price of child services is

$$\Pi(\bar{q}; w, p, P^a) = \frac{\pi^Z(\bar{q}; w, p)}{P^a} = \frac{\mathrm{CCI}(w, p)}{P^a \bar{q}^\eta}. \tag{D.15}$$

Thus, up to the constant $\bar{q}^{-\eta}$, the model's relative price is CCI/P^a.

To see how this relative price enters the household's allocation problem, begin with the fixed-quality benchmark. If quality is held at $\bar{q}$, then $Z = n\bar{q}^\eta$, and the household budget can be written as

$$P^a c + \pi^Z Z = I. \tag{D.16}$$

Dividing by P^a gives

$$c + \Pi Z = \frac{I}{P^a}. \tag{D.17}$$

After expressing resources in units of adult consumption, the only price that governs the tradeoff between adult consumption and child services is Π.

With CES utility,

$$U(c, Z) = \left[(1-\omega)c^{\rho} + \omega Z^{\rho}\right]^{1/\rho}, \qquad \rho = \frac{\sigma - 1}{\sigma}, \tag{D.18}$$

the expenditure share devoted to child services is

$$s_Z(\Pi) = \frac{\omega^{\sigma}\Pi^{1-\sigma}}{(1-\omega)^{\sigma} + \omega^{\sigma}\Pi^{1-\sigma}}. \tag{D.19}$$

This share satisfies

$$\frac{\partial \ln s_Z}{\partial \ln \Pi} = (1-\sigma)(1-s_Z). \tag{D.20}$$

When child services and adult consumption are substitutes, $\sigma > 1$, an increase in the relative price of child services lowers the share of household resources allocated to children.

In the fixed-quality benchmark, fertility is child-services demand divided by the services produced by each child:

$$n^*(\bar{q}) = \frac{Z^*}{\bar{q}^{\eta}} = \frac{s_Z(\Pi)}{\Pi}\frac{I/P^a}{\bar{q}^{\eta}}. \tag{D.21}$$

Holding real full income fixed, changes in fertility satisfy

$$d\ln n^*(\bar{q}) = \left[(1-\sigma)(1-s_Z) - 1\right] d\ln \Pi. \tag{D.22}$$

This expression shows that the fertility decision depends on the cost of child services relative to adult consumption (Π), not the nominal cost of children alone.

The same relative-price logic applies when quality is chosen endogenously. After the cost-minimization step, the household can be written as choosing c, n, and q:

$$\max_{c,n,q} U(c, nq^{\eta}) \qquad \text{s.t.} \qquad P^a c + n\mathcal{C}(q; w, p) = I. \tag{D.23}$$

The first-order conditions for n and q imply

$$\eta\mathcal{C}(q^*; w, p) = q^*\mathcal{C}_q(q^*; w, p) = \pi^q(w, p; A)q^*. \tag{D.24}$$

Using $\mathcal{C}(q; w, p) = wh + \pi^q(w, p; A)q$, this gives

$$q^* = \frac{\eta}{1-\eta}\frac{wh}{\pi^q(w, p; A)}, \qquad \kappa^* \equiv \mathcal{C}(q^*; w, p) = \frac{wh}{1-\eta}. \tag{D.25}$$

The price of one unit of child services at the chosen quality is

$$P_Z^* = \frac{\kappa^*}{(q^*)^\eta}, \qquad \Pi^* = \frac{P_Z^*}{P^a}. \tag{D.26}$$

The household's child-services expenditure share is again $s_Z(\Pi^*)$, and optimal fertility can be written as

$$n^* = \frac{s_Z(\Pi^*)I}{\kappa^*}. \tag{D.27}$$

Log-differentiating gives

$$d\ln n^* = (1-\sigma)(1-s_Z)d\ln\Pi^* + d\ln\left(\frac{I}{P^a}\right) - d\ln\left(\frac{\kappa^*}{P^a}\right). \tag{D.28}$$

Fertility depends on the relative price of child services through the first term, and on real resources per child through the difference between real full income and the real per-child cost.

In the special case with no non-labor income, $V = 0$, full income is $I = w\overline{T}$, while $\kappa^* = wh/(1-\eta)$. The final two terms in equation (D.28) cancel, leaving

$$d\ln n^* = (1-\sigma)(1-s_Z)d\ln\Pi^*. \tag{D.29}$$

Thus, in this benchmark, fertility falls when children become relatively more expensive if child services and adult consumption are substitutes, $\sigma > 1$. More generally, the sign and magnitude of the response depend on preferences, income effects, and the way quality adjusts.

D.4 Valuing forgone leisure

The previous subsections treat parental time as an input into child costs. Empirically, we distinguish between childcare time that comes out of leisure and that which comes out of work. We motivate that here.

Begin with the standard labor-leisure benchmark. Let the household have utility over adult consumption C_t and leisure time L_t, with time constraint

$$H_t + L_t = \overline{T},$$

where H_t is market work. The household faces the budget constraint

$$p^{a_t}C_t = w_t H_t + Y_t,$$

where p^{a_t} is the price of adult consumption, w_t is the wage, and Y_t is non-labor income. At an interior optimum, the first-order condition for leisure is

$$p^{a_t} \frac{U_L(C_t, L_t)}{U_C(C_t, L_t)} = w_t. \tag{D.30}$$

The left-hand side is the dollar value of an additional hour of leisure: the marginal utility of that hour, converted into consumption units and then into dollars. Thus, in the standard interior labor-supply model,

$$q_t^L = w_t. \tag{D.31}$$

This provides the natural benchmark for valuing leisure time. If an hour of childcare displaces an hour of leisure, and the household is free to adjust labor supply at the margin, the lost leisure hour has the same marginal value as the wage.

This benchmark may be incomplete if leisure is itself a home-produced activity that combines time with market goods. Following the household-production logic of Becker (1965), suppose effective leisure services are produced using leisure time L_t and leisure goods X_t^L:

$$\widetilde{L}_t = A_t^L L_t^{1-\alpha_L} \left(X_t^L\right)^{\alpha_L}, \qquad \alpha_L \in [0, 1]. \tag{D.32}$$

Let P_t^L denote the price of leisure goods. The consumption-equivalent value of a leisure hour is now

$$q_t^L \equiv p^{a_t} \frac{U_{\widetilde{L}}}{U_C} \frac{\partial \widetilde{L}_t}{\partial L_t}. \tag{D.33}$$

Using equation (D.32),

$$\frac{\partial \widetilde{L}_t}{\partial L_t} = (1 - \alpha_L) A_t^L \left(\frac{X_t^L}{L_t}\right)^{\alpha_L}. \tag{D.34}$$

Thus, holding other determinants fixed, the value of a leisure hour is increasing in the amount of leisure goods used per hour of leisure. If entertainment, communication, travel, or other leisure goods become cheaper relative to adult consumption, households may combine more of those goods with each hour of leisure. In that case, the marginal value of a leisure hour can rise even if the number of leisure hours is unchanged.

We capture this channel with a reduced-form relationship between leisure-goods intensity and the relative price of leisure goods:

$$\frac{X_t^L}{L_t} \propto \left(\frac{\mathrm{ACI}_t}{P_t^L}\right)^{\kappa}, \qquad \kappa \geq 0. \tag{D.35}$$

Combining equations (D.33)–(D.35), and anchoring the baseline value of leisure to a constant

fraction of the wage, gives

$$q_t^L = \phi w_t \left(\frac{\text{ACI}_t}{P_t^L} \right)^{\theta}, \qquad \theta = \alpha_L \kappa. \tag{D.36}$$

The parameter ϕ governs the level of the leisure valuation, while θ governs how strongly the value of leisure responds to changes in the relative price of leisure goods. ACI_t is the Adult Consumption Index and P_t^L is the leisure-goods price index.

Taking logs of equation (D.36) gives the growth rate used in the leisure-augmented CCI:

$$\Delta \ln q_t^L = \Delta \ln w_t + \theta \left(\Delta \ln \text{ACI}_t - \Delta \ln P_t^L \right). \tag{D.37}$$

When $\theta = 0$, the value of leisure grows with the wage, as in the standard labor-leisure model. When $\theta > 0$, a decline in the price of leisure goods relative to adult consumption raises the value of a forgone leisure hour. The case $\theta = 1$ corresponds to a full adjustment for the relative price of leisure goods.

E Unpacking Childcare Spending Shares

One finding in Section 4 is that childcare and nursery school tuition shares are only around 5% and were relatively stable over this period, despite the sharp rise in their price. This appendix investigates this finding in more detail.

First, we note that the child cost index reflects the budget shares for the average child-year, and therefore includes all children from 0 to 17. Most children above five do not have any spending in that category. When children aged 0-4 are separated from children aged 5-17, the share of spending dedicated to childcare and nursery school increases to between 15% and 20%. Figure E.1 shows that childcare is a major expense for young children, and has risen modestly over time for this subset.

Further, looking even within this group of families with young children, many families report no spending on nursery school or childcare. Figure E.2 shows the share of families with at least one child between 0 and 4 without any spending on childcare or nursery school. Indeed, the majority of families report no childcare spending and this share has risen over time. Among families that do report childcare spending, the childcare share of child spending has risen from roughly 24 percent to over 30 percent (Appendix Figure E.3) This pattern is consistent with a 2011 Census report based on the Survey of Income and Program Participation, which found that only 33 percent of families with a child under five regularly used nonrelative childcare, including care provided through Head Start (Laughlin, 2013).

E.1 Household Spending Around the Birth of a First Child

This appendix presents a descriptive exercise that complements the allocation-based child spending measures used in the main text. The main CCI requires an estimate of the spending bundle associated with a child. Our baseline construction obtains that bundle by allocating household expenditures to children using category-specific rules. An alternative way to learn about child-related spending is to examine how household spending changes around the arrival of a first child. We pursue that exercise here, but use it only descriptively. The reason is that the event-study response to a first birth is not the same object as child spending: it combines direct spending on the child with changes in parents' own consumption, changes in total household resources, and, at longer horizons, the arrival of additional children.

The Consumer Expenditure Survey is a repeated cross-section, so we cannot follow the same household before and after childbirth. We therefore construct a pseudo-event study following the method in Kleven (2025). The sample is restricted to married couples. Post-birth event time is assigned using the age of the oldest child: households with a newborn are assigned event time $\tau = 0$, households whose oldest child is age one are assigned event time $\tau = 1$, and so on. To construct pre-birth event times, households observed with a newborn are matched to childless households who are observationally similar but observed at younger ages. The matching variables include gender, education, region, race, income quintile, and an indicator for living in selected high-cost urban areas. The resulting sample contains approximately 72,500 observations.

For each outcome, we estimate

$$Y_{it} = \sum_{\tau \neq -2} \beta_\tau \mathbf{1}\{\tau_{it} = \tau\} + \alpha_{a(i,t)} + \gamma_t + X_{it}'\delta + \varepsilon_{it}, \tag{E.38}$$

where event time $\tau = -2$ is omitted, $\alpha_{a(i,t)}$ are age fixed effects, γ_t are calendar-year fixed effects, and X_{it} includes demographic controls. The figure below reports estimates for total spending in each broad category. Coefficients are scaled by the category's pre-birth mean at event time $\tau = -2$ and reported as percentage changes. Standard errors are clustered at the individual level.

Appendix Figure E.5 shows that household spending changes substantially around the birth of a first child. Total spending rises after childbirth and remains elevated through much of childhood. The increase is broad-based, but the timing and magnitude differ across categories. Food spending rises steadily after the birth. Although the figure reports only the broad food category, the finer underlying estimates indicate that this aggregate increase combines a rise in food at home with a decline in food away from home, consistent with

children changing both direct food needs and parents' own restaurant consumption.

Education and childcare displays the sharpest age-specific pattern. Spending rises steeply in the early years after birth and then declines as children age. Since the broad category combines education and childcare, this early spike is best interpreted as reflecting paid childcare and nursery-school expenses rather than schooling costs more generally. Healthcare also rises around childbirth and remains above its pre-birth level for much of the post-birth period, consistent with both birth-related health expenses and ongoing medical spending for children. Apparel rises as well, especially around the early childhood years.

Other categories are more useful for illustrating why the event-study coefficients cannot be interpreted as child spending shares. Housing, transportation, communication, and miscellaneous spending all move around the birth of a child, but these responses reflect a mixture of direct child needs, household adjustment, and parents' own consumption responses. For example, a change in housing spending may reflect anticipatory moves, family-size adjustments, or changes in total resources, not simply the amount spent "on" the child. Similarly, changes in communication or miscellaneous spending may reflect parents reallocating their own consumption rather than spending directly on children.

The event-study estimates therefore support the broad premise that children change the composition of household spending, but they do not identify the child-specific spending bundle required for the CCI.

E.2 Child Earnings and Hours Penalty Calculation

Data and sample. We combine the CPS ASEC (1989–2024) and the ACS (2000–2024). Both surveys report individual annual wage and salary income for the previous calendar year, so we date all estimates by income year, $t = 1988, \ldots, 2023$. Hours are usual weekly hours worked (in the CPS, usual hours last year; in the ACS, usual hours in the past 12 months). We convert weekly hours to annual hours as weekly hours $\times$ 50 wherever annualization is needed. For survey years covered by both surveys, we include a survey fixed effect. Survey weights are normalized to sum to one within source $\times$ year, so the CPS and ACS receive equal total influence where they overlap.

For our main estimation, we use the sample of married women and men, aged 20–45, in one of two groups: parents whose oldest own child in the household is aged 0–6, and married individuals without children. We limit to couples whose oldest child is under age 6, and where the parent is aged 20–45, because this is the sample for which childless married women of the same age provide the best earnings comparison. At older child ages, the contemporaneous-childless comparison becomes increasingly selected, biasing the estimated

penalties (see Appendix Figure A.10). We discuss how we compute the penalty for parents of children aged 7+ below.

Estimation and Construction of Penalty. We estimate equation (17) on the above sample. The coefficient $\beta_{a,t}$ is the level penalty: the earnings (or hours) gap between a mother whose oldest child is age a and an observationally identical childless married woman in survey year t. To express penalties in percent terms, we divide each level penalty by a counterfactual: the earnings a mother of a child of age a would have in year t absent children. We construct this as a weighted average of the earnings of childless married women in year t, where the weights are the 1990 distribution of mothers' ages among children of age a.[26] Weighting by the 1990 age distribution of mothers ensures that changes in the counterfactual reflect changes in childless women's earnings as opposed to changes in the age composition of mothers (e.g., delayed fertility). The percent penalty is then $\pi_{a,t} = \beta_{a,t}/\bar{Y}^{\text{cf}}_{a,t}$.

Because we exclude parents of children aged 7-17 in our main sample, we hold the percent penalty fixed at its age-6 value ($\pi_{a,t} = \pi_{6,t}$) and price it with an age-appropriate counterfactual. Specifically, the mother-age weights for child age $a \geq 7$ are the age-6 weights shifted up by $a - 6$ years (capped at 60), applied to current-year childless women aged 20–60. The level penalty at these ages is then $\beta_{a,t} = \pi_{6,t} \cdot \bar{Y}^{\text{cf}}_{a,t}$. Fixing the penalty in percent rather than dollar terms means it grows with the counterfactual as women age, which we view as the conservative choice; because the true penalty likely declines as children age, these extrapolated values are an upper bound.

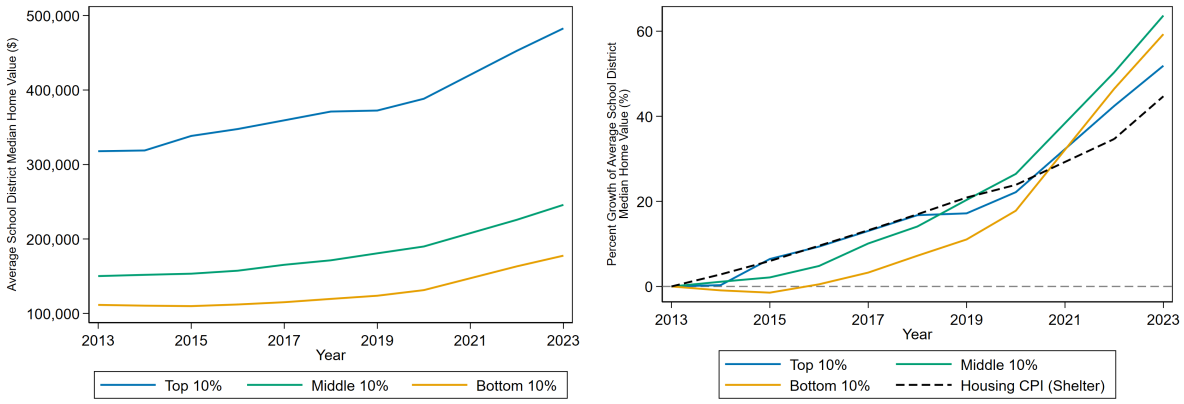
To aggregate across child ages, we first smooth the level penalties and the counterfactuals with a five-year centered moving average and recompute percent penalties from the smoothed series. Then, we average the age-specific penalties within each year using the 1990 distribution of children's ages as weights. Averaging the age-specific products, rather than multiplying an average percent penalty by an average counterfactual, preserves the covariance between penalties and counterfactuals across child ages.

An alternative way of estimating the child earnings and hours penalties is to follow Kleven (2025) and estimate event studies by matching women with a newborn with observationally similar women 1, 2, 3, 4, or 5 years younger. While the penalty estimates are similar (see Figure E.6), the counterfactual earnings used in our calculation are disproportionately drawn from younger women.

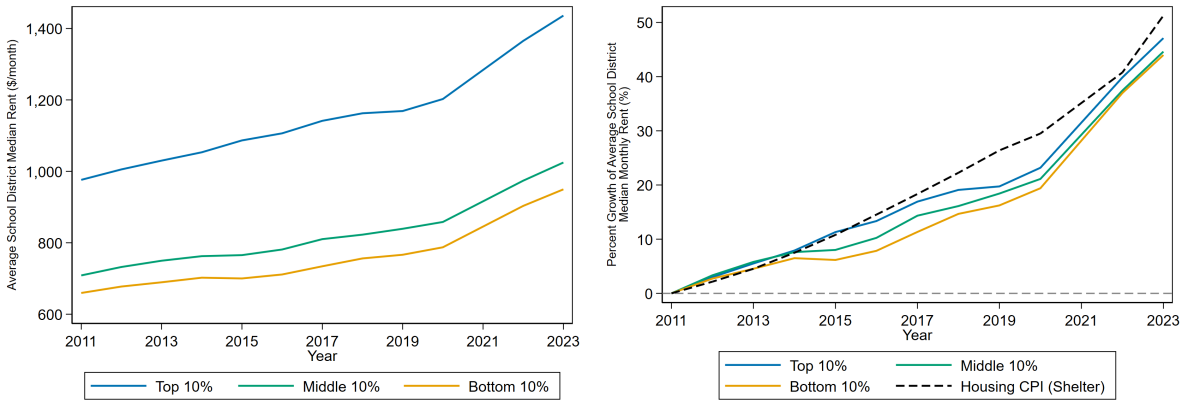
[26] The ages are combined as 1-2, 3-4, and 5-6.

Figure C.1: Housing-cost growth by school quality

(a) **Owner-Occupied Home Value** (b) **Owner-Occupied Home Value Growth**



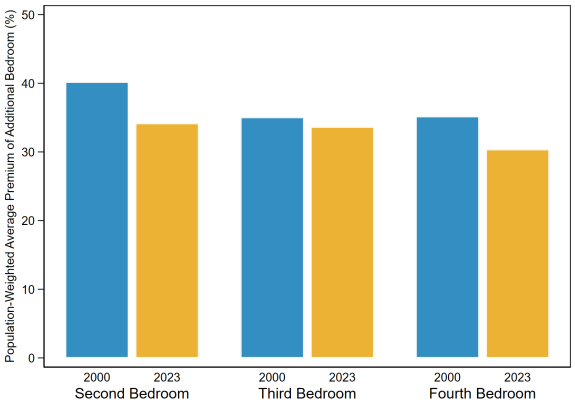
(c) **Monthly Rent** (d) **Monthly Rent Growth**



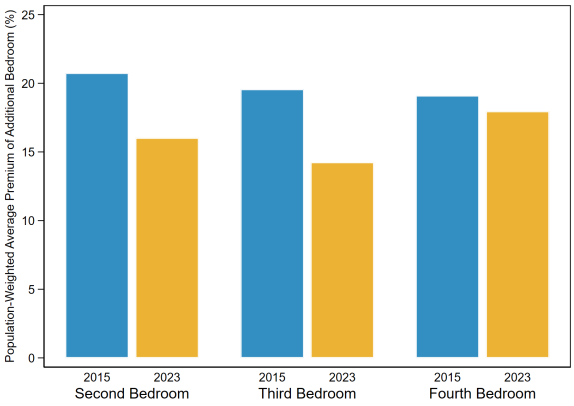
Notes: Panels (a) and (c) plot the average school-district median owner-occupied home value and median gross rent, respectively, for districts in the top 10 percent, middle 10 percent (45th–55th percentiles), and bottom 10 percent of the test-score distribution within their state. Districts are re-ranked in each year using test scores pooled over the same trailing five-year window as the ACS outcome. Panels (b) and (d) plot cumulative percent growth since the first sample year (2013 for owner values and 2011 for rents). The black dashed line reports cumulative growth in the BLS shelter CPI.

Figure C.2: Differential housing-cost growth by bedroom count

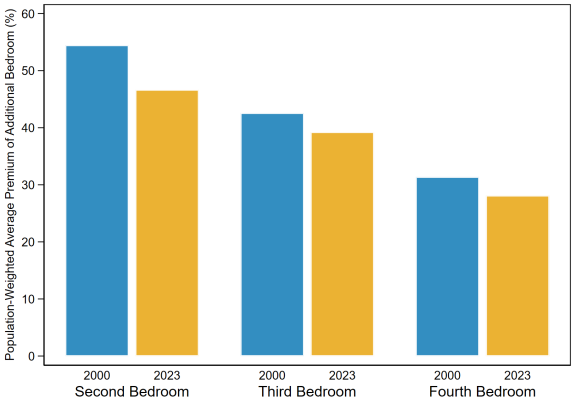
(a) **City comparison: Home Values**



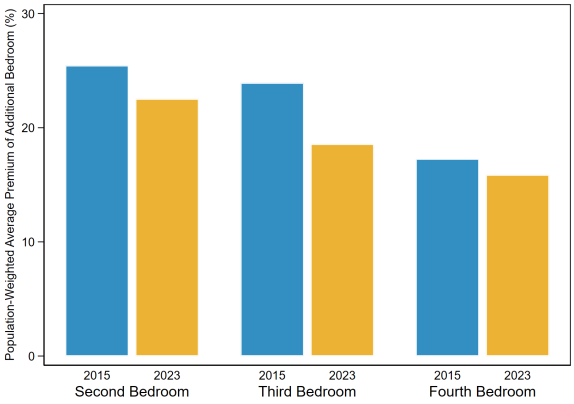
(b) **State comparison: Rents**



(c) **ZIP comparison: Home Values**



(d) **County comparison: Rents**



Notes: Each pair of bars reports the population-weighted average bedroom premium in the indicated endpoint years. Panels (a) and (c) use the Zillow Home Value Index for cities and ZIP codes and compare 2000 with 2023. Panels (b) and (d) use ACS median gross rent for states and counties and compare 2015 with 2023. Each bedroom transition uses the pairwise sample of geographies reporting both adjacent bedroom categories in both endpoint years. A lower 2023 bar means that the larger unit appreciated less rapidly than the adjacent smaller unit.

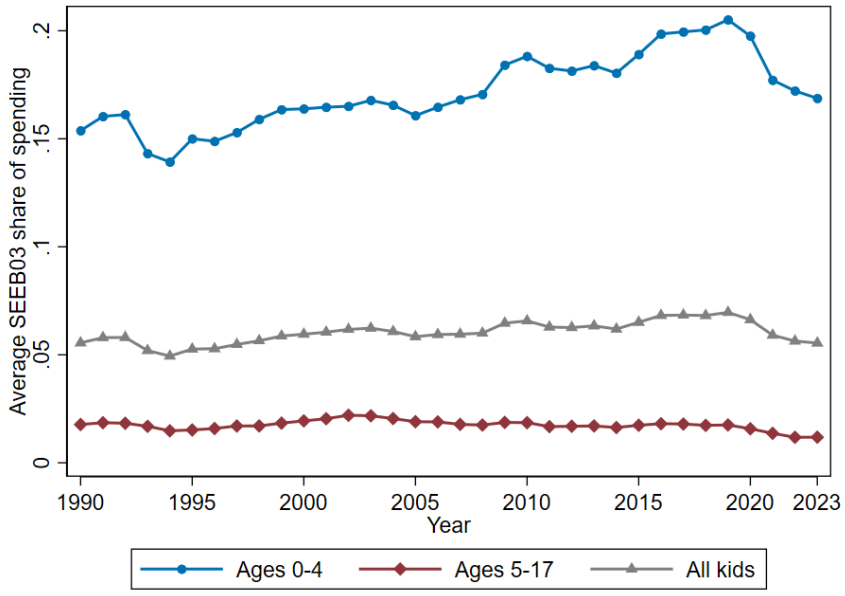


Figure E.1: Childcare/Nursery School Spending by Age Group

Notes: This figure shows the share of children's spending each year that is allocated to childcare and nursery school tuition. The grey line depicts the share of spending for all children aged 0-17, weighting each age equally. The red line is the share of spending for children aged 5-17, with each age's share equally weighted. The blue line is the equally age weighted childcare spending share for children aged 0-4.

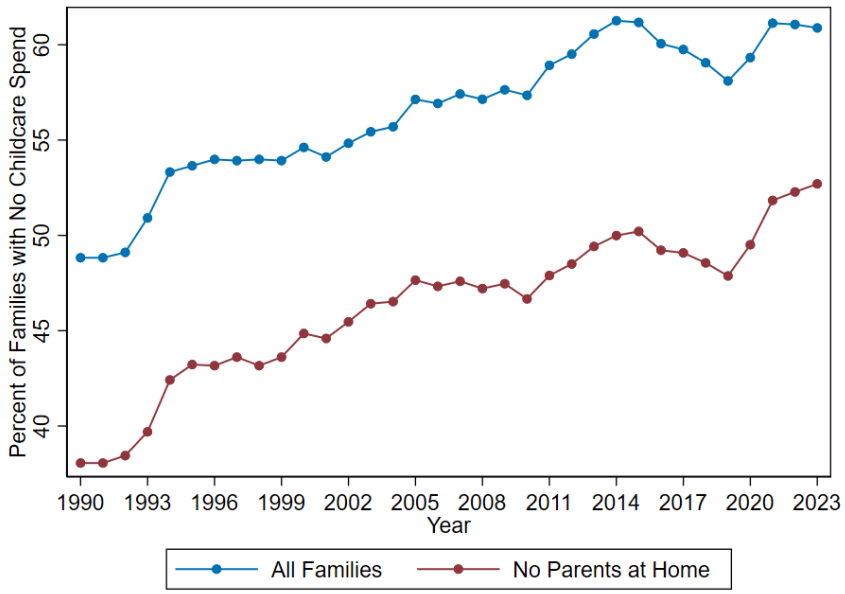


Figure E.2: Share of Families with a Child 0-4 and No Childcare Expenses

Notes: This figure shows the share of families that have at least one child under 5 that report no spending on childcare or nursery school tuition. This figure uses CEX data that pools up to five years of data together for each target year. The blue line is all families with at least one child under 5. The red line is all families with at least one child under five where all parents work.

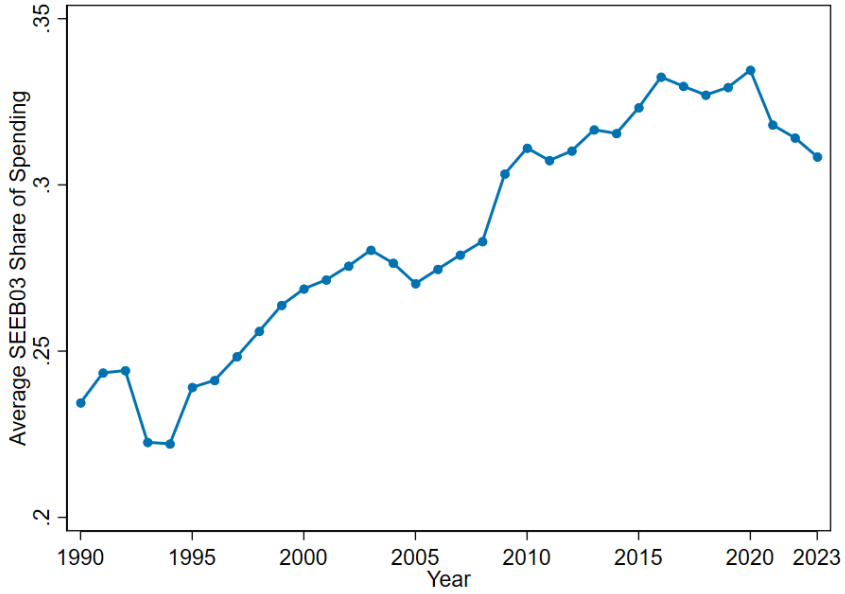


Figure E.3: Share of Spending on Childcare for Children 0-4, Given > 0 Childcare Spending

Notes: This figure shows the share of spending that is dedicated to childcare and nursery school tuition for children who are aged 0-4 and have positive spending in that category.

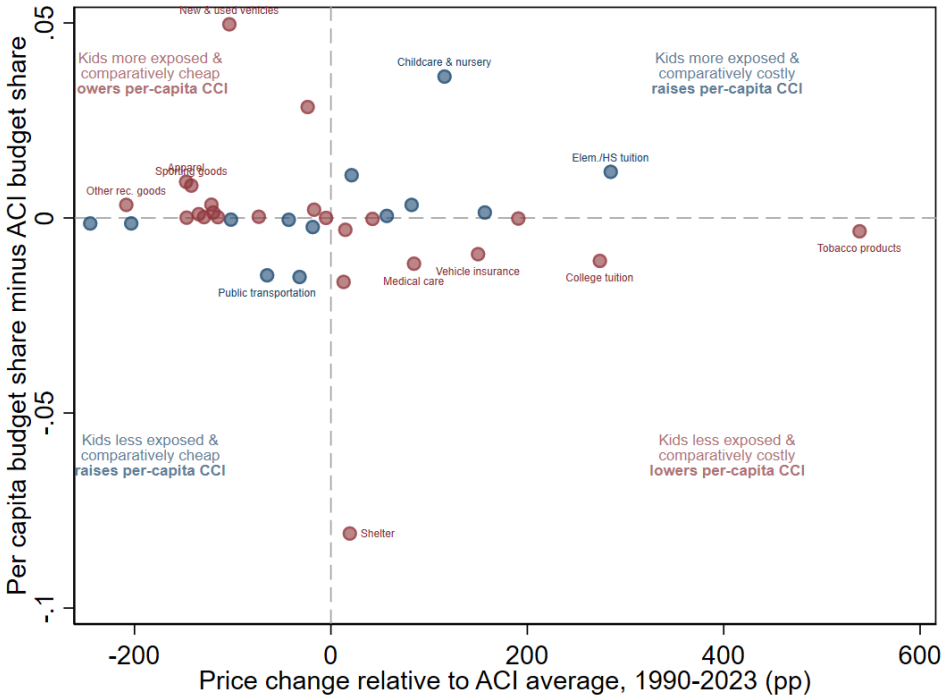
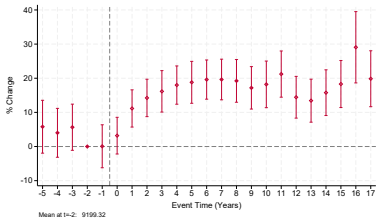


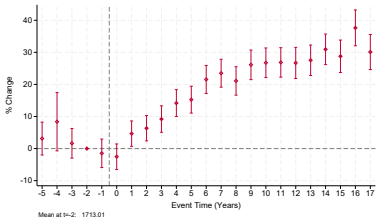
Figure E.4: Category contributions to the child–adult inflation gap: Average ACI

Notes: Each marker corresponds to a spending category. The horizontal axis is cumulative price growth in that category relative to the ACI from 1990 to 2023. The vertical axis is the category's average budget share in the child basket minus its average budget share in the adult benchmark. Marker area is proportional to the category's average budget share in the child basket. Navy markers make positive contributions to CCI^g relative to the ACI; maroon markers make negative contributions. Shares are constructed from the CEX and prices from BLS CPI subindices.

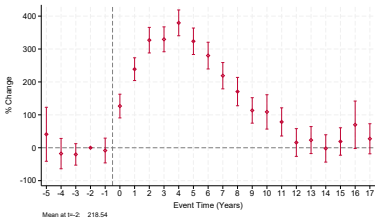
Figure E.5: Household spending around the birth of a first child



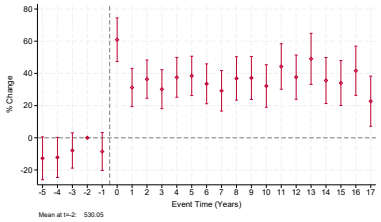
(a) Total spending



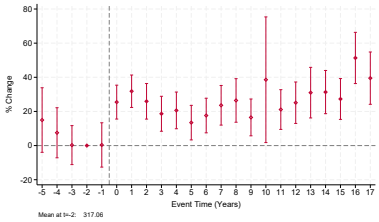
(b) Food



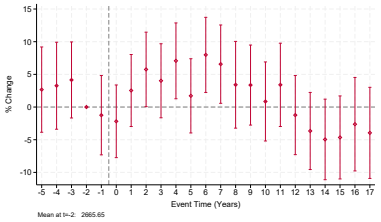
(c) Education and childcare



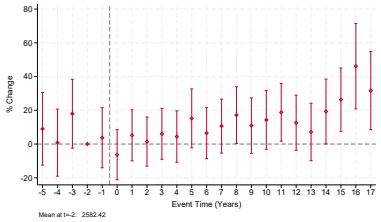
(d) Healthcare



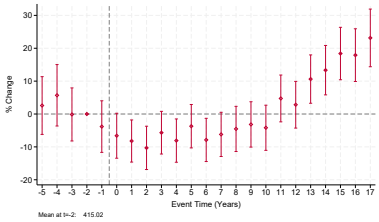
(e) Apparel



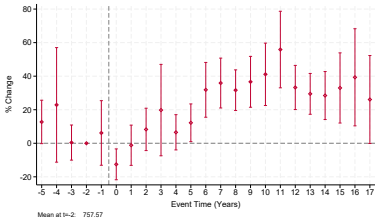
(f) Housing



(g) Transportation



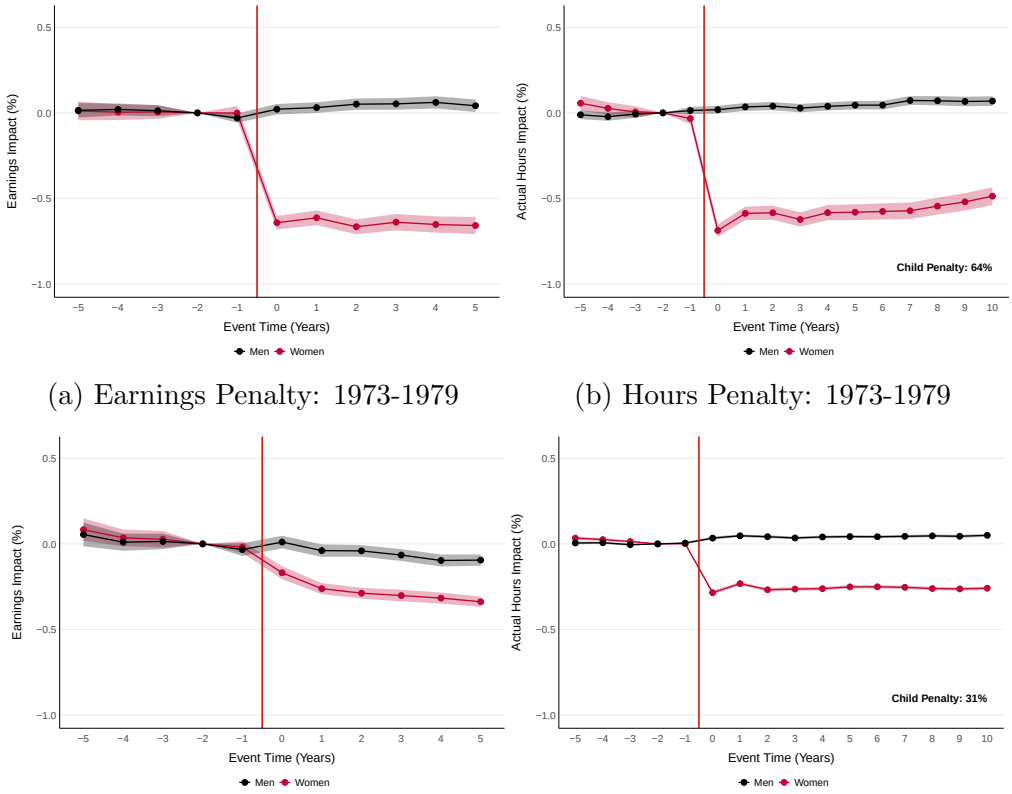
(h) Communication



(i) Miscellaneous

Notes: The figure reports pseudo-event-study estimates of household spending around the birth of a first child. Outcomes are total spending in each broad category, with coefficients converted to percent changes relative to the category mean at event time $\tau = -2$. Event time $\tau = -2$ is omitted. The sample is restricted to married couples. Post-birth event time is assigned using the age of the oldest child; pre-birth event times are constructed by matching households with a newborn to observationally similar childless households observed at younger ages. Estimates include event-time indicators, age fixed effects, calendar-year fixed effects, and demographic controls. Standard errors are clustered at the individual level.

Figure E.6: Child Penalty Estimates following Kleven (2025)



(a) Earnings Penalty: 1973-1979

(b) Hours Penalty: 1973-1979

(c) Earnings Penalty: 2020-2025

(d) Hours Penalty: 2020-2025

Note: This figure shows five- and ten-year earnings and hours penalties following the methodology in Kleven (2025).

E.3 Home Production and Parental Time

The time budget of an individual with children must satisfy the identity

$$\Delta \mathrm{CC}_t \;=\; |\Delta W_t| + |\Delta L_t| - \Delta \mathrm{HP}_t,$$

where $\Delta \mathrm{CC}_t$ is the increase in childcare hours, $|\Delta W_t|$ is the reduction in market work, $|\Delta L_t|$ is the reduction in leisure, and $\Delta \mathrm{HP}_t$ is the change in home production (positive means more home production). Because children generate additional home production—more meals to prepare, more laundry, more household tasks—$\Delta \mathrm{HP}_t > 0$ in the data, so $|\Delta W_t| + |\Delta L_t| > \Delta \mathrm{CC}_t$.

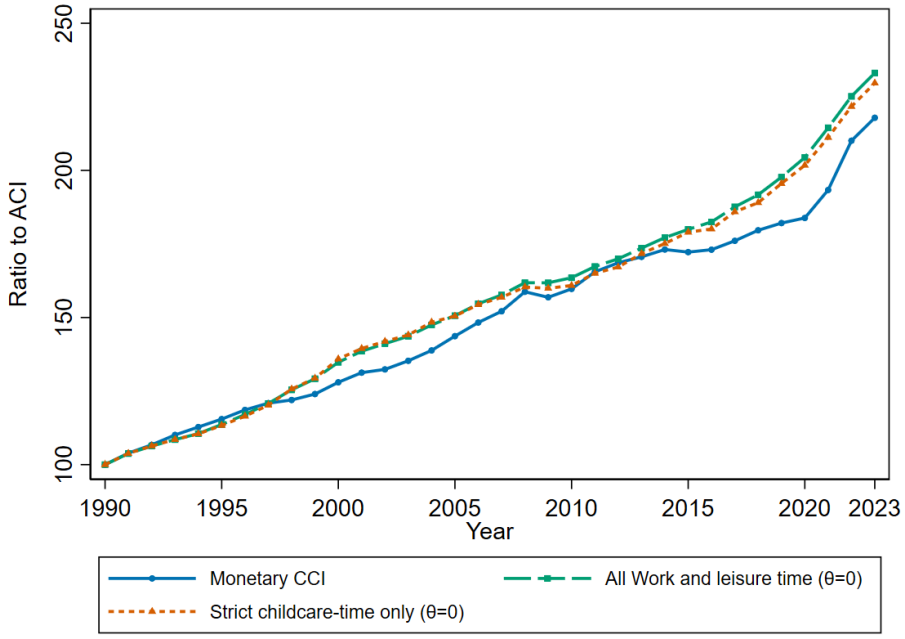
In our baseline index, the time component of the CCI prices the full reductions in market work and leisure. Because the sum of these two exceeds measured childcare hours for women, the baseline implicitly assigns a cost to the incremental home production induced by children, treating it as an additional margin through which children claim parental time. This is the appropriate treatment if child-related home production represents a genuine use of time that parents would not otherwise perform.

An alternative view is that home production induced by children is partly an endogenous increase in household production for the parents as well. For example, if the woman reduces work hours to be home with the children, she has more time to mow the lawn, do adult laundry, cook family dinners, etc. Under this view, the increase in home production should not be counted as an input in the CCI. We implement this alternative by re-weighting the work and leisure hour inputs so that they sum to the measured childcare increment $\Delta \mathrm{CC}_t$, scaling each proportionally:

$$\widetilde{e}_t^M = e_t^M \cdot \frac{\Delta \mathrm{CC}_t}{e_t^M + e_t^L}, \qquad \widetilde{e}_t^L = e_t^L \cdot \frac{\Delta \mathrm{CC}_t}{e_t^M + e_t^L}.$$

In practice, the two approaches yield very similar results (See Figure E.7). We therefore adopt the baseline (full work and leisure reduction), and treat the childcare-normalized version as a robustness check.

Figure E.7: Robustness: Accounting for home-production



Notes: Each series shows a price index, normalized to 100 in 1990. The monetary CCI includes only children's goods and services. The "All work and leisure time" series prices all childcare hours drawn from market work (at the effective wage w_t^{eff}) and leisure (at the shadow wage). The "Strict childcare-time only" series restricts the priced time to the measured childcare increment only, scaling work and leisure hours proportionally so that they sum to the childcare time increase rather than the total work-plus-leisure reduction. For each treatment, $\theta = 0$ values leisure at the wage.