



The Budget Lab's Early Childhood Education and Care Model

Published: April 16, 2026

Introduction

This report describes the Budget Lab's model for evaluating the impact of policies affecting early childhood education and care (ECEC). The model estimates how federal policy initiatives, like new subsidies for childcare, would affect:

- **patterns of childcare usage and parental employment**, which are closely related decisions for families;
- **cost burdens for different types of families**, accounting for changes in subsidy eligibility and market prices;
- **the prices that childcare providers charge**;
- **the wages earned by care workers**;
- **the federal budget**, including feedback effects on tax revenues from changes in parental employment; and
- **children's future economic outcomes**, based on how a policy reform induces children to move between ECEC options with different estimated impacts on their future earnings.

To capture these features of ECEC policy, we build on the work of [Borowsky et al. \(2022\)](#) and Brown et al. (2026)¹ and develop a structural equilibrium microsimulation model of parents' joint employment and care decision.² Family-level decisions aggregate to market demand for different forms of ECEC (for example, center-based care vs. home-based care), and provider prices and wages adjust to clear the market. Policy reforms can shift the equilibrium, and the model calculates how those changes flow through to fiscal and distributional effects.

In what follows, we describe the structural model setup, the data sources and associated statistical modeling, the calibration process, and more.

Equilibrium Model

Notation

In this section, *italic* variables indicate scalars, while **bold** variables indicate vectors. A bar over a variable (e.g. \bar{w}) indicates an exogenous quantity, while a hat (e.g. $\hat{\theta}$) indicates a quantity derived from estimation or calibration.

Childcare sectors are denoted by s . Time is denoted by t .

Policy

A policy rule $\theta_t : \mathbb{R}^S \rightarrow \mathbb{R}^S$ is a function that maps market prices \mathbf{p}_t to prices faced by families and providers, with the former elements of $\theta_t(\mathbf{p}_t)$ denoted $\theta_t^D(\mathbf{p}_t; \bar{\mathbf{d}}_t)$ (with $\bar{\mathbf{d}}_t$ defined below) and the latter denoted $\theta_t^S(\mathbf{p}_t)$ (with typical element $\theta_t^s(p_t^s)$ denoting the post-policy price faced by firms in sector s). In our baseline, $\theta_t(\mathbf{p}_t) = \mathbf{p}_t$ for all t .

Demand

In our model, as in [Borowsky et al. \(2022\)](#), families make joint decisions over maternal labor supply and ECEC arrangements. Each family (parent unit) i , with N_i children, chooses among $j = 1, 2, \dots, J$ options. Each option j is an $(2N_i + 1)$ -tuple consisting of:

- The childcare sector s_i^n for each child (where $n \in \{1, \dots, N_i\}$ and $s_i^n = s_0$ denotes parental care);
- Hours of non-parental childcare $h_i^n \in \{\text{None, Part-Time, Full-Time}\}$ for each child; and
- A maternal labor-supply choice $l_i \in \{\text{Non-Employed, Employed Part-Time, Employed Full-Time}\}$.

At time t , the household chooses the option j that maximizes utility

$$U_{ijt} = V_{ijt} + \varepsilon_{ij}$$

where

$$V_{ijt} = \alpha_{ij} + \beta \cdot \left(\frac{(Y_{ijt} - T_{ijt} - C_{ijt})^{1-\rho}}{1-\rho} \right) + \Delta_t \cdot 1_{j \in \text{Emp}}$$

Here:

- Y_{ij} is i 's pretax income (AGI) associated with option j ;
- T_{ij} is i 's total tax liability associated with option j ;
- C_{ij} is i 's total cost of childcare associated with option j (the product of the after-policy price and care hours associated with option j);
- $\alpha_{ij} \in \mathbb{R}$ is an individual/option-specific intercept term; and
- $\beta \in \mathbb{R}_+$ and $\rho \in [0, 1)$ are utility parameters common to all parent units; and
- Δ_t is a time-varying utility shifter for options involving part- or full-time maternal employment (we use this shifter to benchmark our simulated baseline employment rates by age and sex in future years to projections from the Congressional Budget Office [CBO]).

Assuming that ε_{ij} are (time-invariant) i.i.d. Type-I extreme value, the probability that parent unit i chooses option j at time t is thus

$$P_{ijt} = \frac{\exp(V_{ijt})}{\sum_m \exp(V_{imt})}$$

These choice probabilities can be aggregated across parent units to generate an aggregate demand function $D(\mathbf{p}; \mathbf{d}, \Psi)$, where \mathbf{p} are prices, \mathbf{d} are exogenous determinants of demand (demographics, potential wages [described below], etc.), and Ψ are (time-invariant) parameters of the demand function (α_{ij} , β and ρ).

The vector of quantities demanded at time t given our calibrated parameters of the demand function and exogenous determinants of demand can be written as:

$$\mathbf{Q}_t^D = \mathbf{D} \left(\theta_t^D (\mathbf{p}_t; \bar{\mathbf{d}}_t); \bar{\mathbf{d}}_t, \hat{\Psi} \right)$$

Supply

The supply side of our model is taken from [Borowsky et al. \(2022\)](#). ECEC services in sector s are produced with a Leontief (fixed-proportions) production technology:

$$Q^s = \min\{a^{BA,s} l^{BA,s}, a^{NoBA,s} l^{NoBA,s}\},$$

where $l^{BA,s}$ and $l^{NoBA,s}$ are hours of labor supplied by care workers with and without a bachelor's degree, respectively.

Denote care worker wages (which are assumed to be the same across all S sectors) as $\mathbf{w}_t \equiv (w_t^{NoBA}, w_t^{BA})$, and denote the estimated parameters of the supply function by $\hat{\mathbf{a}} \equiv \{\hat{a}^{NoBA,s}, \hat{a}^{BA,s}\}_s$. The vector of quantities supplied at time t can then be written as:

$$\mathbf{Q}_t^S (\theta_t^S (\mathbf{p}_t), \mathbf{w}_t; \bar{\delta}_t, \hat{\mathbf{a}}) = (Q_t^1, \dots, Q_t^s, \dots, Q_t^S)$$

We specify wages for low- and high-skill childcare workers via a constant-elasticity sectoral labor supply function. Let $\bar{\mathbf{w}}_t = (\bar{w}_t^{NoBA}, \bar{w}_t^{BA})$ be the vector of nominal wages implied by the aggregate nominal wage growth series in year t (see the Initialization step below). Then equilibrium childcare wages are given by

$$\mathbf{w}_t (\mathbf{Q}_t^S) = (w_t^{NoBA} (\mathbf{Q}_t^S), w_t^{BA} (\mathbf{Q}_t^S)) = \left(\bar{w}_t^{NoBA} \cdot \left(\frac{\sum_s \left(\frac{Q_t^s}{\hat{a}^{NoBA,s}} \right)}{\sum_s \left(\frac{Q_t^s}{\hat{a}^{NoBA,s}} \right)} \right)^{\frac{1}{\bar{\eta}^{NoBA}}}, \bar{w}_t^{BA} \cdot \left(\frac{\sum_s \left(\frac{Q_t^s}{\hat{a}^{BA,s}} \right)}{\sum_s \left(\frac{Q_t^s}{\hat{a}^{BA,s}} \right)} \right)^{\frac{1}{\bar{\eta}^{BA}}} \right)$$

where $\bar{\eta}^{NoBA}$ and $\bar{\eta}^{BA}$ are time-invariant elasticities of childcare labor supply for workers without and with a bachelor's degree, respectively.

As [Borowsky et al. \(2022\)](#) note, there are only a limited number of studies that estimate the labor-supply elasticities of potential care workers. Following their methodology, we choose parameter values of $\bar{\eta}^{NoBA} = 4$ and $\bar{\eta}^{BA} = 2$, which are on the low end of the range of elasticities in the literature. Since we err on the side of care workers' labor supply being more inelastic, our estimated effects on wages represent something of an upper bound.

The quantity supplied by sector s is given implicitly by the *per-unit residual revenue equation*:

$$\bar{\delta}_t^s = \theta_t^s(p_t^s) - \left(\frac{w_t^{NoBA}(\mathbf{Q}_t^S)}{\hat{a}^{NoBA,s}} \right) - \left(\frac{w_t^{BA}(\mathbf{Q}_t^S)}{\hat{a}^{BA,s}} \right)$$

Note that the vector of residual per-unit revenues $\bar{\delta}_t \equiv (\bar{\delta}_t^1, \dots, \bar{\delta}_t^s, \dots, \bar{\delta}_t^S)$ is assumed to be exogenous and fixed - as in Gibbs et al. (forthcoming) and Borowsky et al. (2022), we take residual revenue as a per-unit quantity that includes nonlabor costs and profits (for for-profit providers).

Equilibrium Algorithms

Given a policy function θ_t , the equilibrium vector of prices \mathbf{p}_t^* and quantities \mathbf{Q}_t^* satisfy the numbered equations above along with the market-clearing condition

$$\mathbf{Q}_t^* = \mathbf{Q}_t^D = \mathbf{Q}_t^S$$

Initialization

At time t , compute all exogenous variables as follows:

- **Per-unit residual revenue ($\bar{\delta}_t$):** Multiply $\bar{\delta}_{t-1}$ by Consumer Price Index (CPI) inflation.
- **Wage index (\bar{w}_t):** Multiply \bar{w}_{t-1} by economy-wide nominal wage growth per hour worked.
- **Demand inputs ($\bar{\mathbf{d}}_t$):** Update weights using CBO demographic projections and multiply potential wages by economy-wide nominal wage growth. We then deflate by CPI indexed to 2019 such that inflation does not affect choices over time.

Equilibrium Search

For a given year t and policy function θ_t , we compute the equilibrium vector of prices and quantities $(\mathbf{p}_t^*, \mathbf{Q}_t^*)$ as follows:

1. Choose an initial price vector \mathbf{p}_t (in practice, we set $\mathbf{p}_t = \mathbf{p}_{t-1}^{*,base}$ [for baseline] or $\mathbf{p}_t = \mathbf{p}_t^{*,base}$ [for scenario]).
2. Compute post-policy prices $\theta_t^D(\mathbf{p}_t; \bar{\mathbf{d}}_t)$ for households.
 1. If finding the equilibrium for our baseline, find Δ_t so that maternal employment matches CBO projections.

Compute aggregate demand $\mathbf{Q}_t^D = \mathbf{D}(\theta_t^D(\mathbf{p}_t; \bar{\mathbf{d}}_t); \bar{\mathbf{d}}_t, \hat{\Psi})$.
3. Using the quantities demanded for each sector (Q_t^s) computed in the previous step, compute wages $\mathbf{w}_t(\mathbf{Q}_t^D)$ using equation (2) above.
4. Compute post-policy prices $\theta_t^S(\mathbf{p}_t)$ for sectors, then, using the wages \mathbf{w}_t computed in the previous step, compute per-unit residual revenue $\tilde{\delta}_t^s$ using equation (3) for each sector. Collect these into $\tilde{\delta}_t \equiv (\tilde{\delta}_t^1, \dots, \tilde{\delta}_t^S)$.
5. Compare $\tilde{\delta}_t$ and $\bar{\delta}_t$ and iterate until they are equal (within some tolerance).

(Note that this algorithm guarantees that the market-clearing condition (4) holds in equilibrium.)

Data

We combine two main data sources to construct the model's underlying household microdata:

- **The American Community Survey (ACS)**, a large and nationally representative survey with high-quality demographic and economic information.
- **The National Survey of Early Care and Education (NSECE)**, a set of smaller ECEC-focused surveys of households and providers.

To combine these two datasets, we perform a kind of statistical matching exercise where we estimate patterns of ECEC demand conditional on observable household characteristics, then impute values in the ACS conditional on those characteristics. This section details the steps involved in this process.

NSECE Household Processing

The NSECE household survey collects information about childcare arrangements in a reference week. We process this data to arrive at a single primary care arrangement per child, which we then use to train our imputation models.

The raw NSECE comes as wide-format tabular dataset where each row represents one household and the columns encode complex relationships within and outside of the household (for example, child-parent linkages and child-provider linkages). We transform this data into a hierarchy of six linkable tables: households, household members, parent units, children, providers, and enrollment records. The *parent unit* is the model's core decision-making unit of analysis, representing either a single parent or a married (or cohabitating) couple. Each parent unit may have multiple children under age 5, and there can be multiple parent units in one household.

Some variables in the NSECE data contain missing values, and some records are outside the scope of our analysis. To arrive at our estimation sample, we apply the following filtering conditions:³

- Children must be aged 0 through 4 at the time of the survey, as the model focuses on ECEC prior to kindergarten entry.
- ECEC arrangements (child-provider pairs) must have non-missing values for both type and hours classification.
- ECEC arrangements must have nonzero weekly hours reported to be considered in the primary care selection algorithm described below.
- Parent units cannot have more than two identified parents implied by the relationship variables. Parent units must have non-missing labor supply variables.

Next, we build the decision set for our discrete choice framework. The first step in this process is to classify ECEC arrangements into one of several economically meaningful categories. We follow Borowsky et al. (2022) and classify arrangements as belonging to one of seven mutually exclusive categories:

- **Unpaid center-based care:** Center-based program with zero reported out-of-pocket cost (e.g. a Head Start program or public pre-K).
- **Low-priced center-based care:** Center-based program with positive out-of-pocket cost below median price per hour.
- **High-priced center-based care:** Center-based program with positive out-of-pocket cost at or above median price per hour.

- **Paid home-based care:** Individual care outside the child's home, with no prior relationship with caregiver (i.e. a home-based daycare), with positive out-of-pocket cost.
- **Unpaid home-based care:** Same as above but with zero reported out-of-pocket cost.
- **Other paid care:** In-home care provided by someone with a prior relationship (like a nanny), or organizational/irregular arrangements (like care provided during religious services), with positive out-of-pocket cost.
- **Other unpaid care:** Same as above but with zero reported out-of-pocket cost.

The second step is to aggregate ECEC arrangements up to the child level by defining a notion of primary care. Respondents often report multiple arrangements (for example, a child might attend pre-school for half-days and stay with a grandparent in the afternoon), but our discrete choice framework requires specifying a single, primary arrangement for each child. Again, following Borowsky et al. (2022), we use a priority-based algorithm that generally privileges formal, paid arrangements over informal, non-paid arrangements.⁴ For each child, we filter to "regular" (i.e. used on a recurring weekly basis) arrangements then aggregate hours across care type. Then, we apply the following priority ranking: high-priced center-based care, low-priced center-based care, unpaid center-based care, paid home-based care, other paid care, unpaid home-based care, other unpaid care. Ties are broken by total hours. If a child has no reported ECEC arrangements, they are assigned "parent-only" status.

The final step in NSECE processing is to deal with families with more than two young children. Our model's choice set is designed for families with one or two children under age five; families with three or more young children present computational challenge given combinatoric expansion. We address this by splitting larger families into two-child "pseudofamilies" representing all unordered pairs of children.

A family with m children generates $\binom{m}{2}$ pseudofamilies. For a three-child family, this produces three pseudofamilies: children (1, 2), children (1, 3), and children (2, 3). Each pseudofamily is treated as an independent observation in the simulation. Family-level weights (used for parent and household counts) are divided by $\binom{m}{2}$. This ensures that the original family is counted exactly once when summing across pseudofamilies. Child-level weights (used for child counts) are divided by $(m - 1)$ because each child appears in exactly $(m - 1)$ pseudofamily pairs.

ACS Processing and Imputations

The American Community Survey provides a large, nationally representative sample with high-quality demographic and economic information, but it lacks detail on ECEC usage and costs. To construct the simulation dataset, we process the ACS microdata and statistically impute ECEC enrollment patterns and associated costs using models trained on the NSECE household survey.

The first step in our processing is to transform the raw ACS dataset, which we download from IPUMS. It is structured as person-level records with household identifiers and relationship pointers, into the same hierarchical format of linkable tables mirroring the NSECE processing output: households, household members, parent units, children, and tax units. Household and tax unit concepts come with the data; parent units need to be constructed.

To build parent units, we start by filtering to children under age thirteen. (While the simulation focuses on the ECEC arrangements of children under age five, we still track the presence of older children in each parent unit, since these children affect tax liabilities and labor supply decisions.) For each child, we identify up to two linked

parents within the household. Children are then grouped by their unique parent identifier to form parent units. This process results in parent units having either one parent or two parents, the latter case encompassing both married and unmarried but cohabitating parents. In one-parent cases, "primary caregiver" status—the person for whom employment status can change with policy—is assigned to the only parent; in two-parent cases, the primary caregiver is assigned to be the mother for opposite-sex couples and assigned arbitrarily for same-sex couples. We also apply the same "pseudofamily" splitting algorithm described in the NSECE section above to split families with 3+ children under age five into several records containing exactly two children under age five.

The next step is to impute latent per-hour labor earnings (including both wages and self-employment income) for ACS records where the primary caregiver does not work. This is required so that we can simulate net income and thus utility under counterfactual employment scenarios in the simulation. To address selection bias, we use a Heckman sample selection model consistent of two equations estimated jointly via maximum likelihood:

1. **Selection equation (employment).** This is a probit model estimating the probability that the primary caregiver is employed as a function of age (quadratic), sex, race, Hispanic origin, educational attainment. To identify the model, in this stage we include as covariate the number of children in the parent unit aged 0--2, 3--5, and 6--12. These variables are assumed to affect labor supply but not wage offers.
2. **Outcome equation (earnings).** This is a linear regression model of the log of hourly earnings for employed, expressed as a function of age (quadratic), sex, race, Hispanic origin, and educational attainment.

We estimate these parameters separately for one- and two-parent cases. After fitting the models, we predict unconditional (latent) hourly earnings for all parents, including non-workers. To preserve realistic variation, we add stochastic noise to predictions by drawing from a normal distribution with mean zero and standard deviation equal to the model's residual standard error. For primary caregivers with observed positive earnings, we retain actual values; for non-workers, we use the stochastic prediction.

Next, we move to ECEC arrangement imputation. The goal is to estimate the distribution of ECEC arrangements conditional on observable parent unit characteristics and then assign those conditional distributions to records in the ACS data. To do this, we train random forest models on the processed NSECE enrollment data and apply them to predict enrollment probabilities in the ACS.

The models predict the joint ECEC arrangement for all children in a parent unit. For one-child families, the outcome variable is a single categorical outcome representing the child's ECEC type (one of eight types) and intensity (part-time or full-time, or NA in the case of parent-only care), with 15 unique options. For two-child families, the model predicts a joint outcome capturing both children's arrangements and intensities simultaneously (e.g., "child 1 in full-time center care, child 2 in part-time home-based care"), preserving within-family correlation in care choices. There are $15^2 = 225$ options in this case.

The model considers the following conditioning variables: age and sex of each child, number of parents, age of each parent, employment status of primary caregiver, sex of primary caregiver, household income percentile, household size, and Census region.⁵ All variables are considered at each decision tree split.⁶

The random forest returns a probability distribution over all possible ECEC arrangements for a given employment status. To generate the full joint distribution of care and employment choices, for each record, we fit values under each of the three cases for employment status, then multiply through the unconditional probability of each employment status (the estimation of which is discussed later in this section).⁷

After imputing ECEC choice probabilities, the next step is to impute values reflecting unmodeled price heterogeneity in out-of-pocket costs by ECEC type. Even within the same ECEC sector (e.g. high-priced center-based care), families face considerable price variation due to both policy factors (i.e. subsidies) and economic factors (e.g. quality differences and geographic variation). Modeling this variation in our simulation (as opposed to simply assuming each family faces the national mean price) is important for accurate calculation of net income under nonlinear policy design.⁸

To capture this heterogeneity, we estimate household-specific price multipliers ("price wedges") that map from supply-side market prices to household out-of-pocket costs. We define the price wedge as the ratio of out-of-pocket hourly cost (reported by families in the NSECE) to the sector-specific supply price (computed from the NSECE provider surveys). A price wedge less than (greater than) 1 indicates the household pays below (above) the national average market rate. Importantly, the NSECE household-reported payment variable already reflects any subsidies received in the baseline policy environment. We do not separately model baseline subsidies; instead, the price wedge absorbs all sources of price variation, including baseline subsidy receipt. Policy reforms in the simulation introduce new subsidies on top of this baseline heterogeneity.

To impute these price wedges, we train separate quantile regression forests for each paid ECEC sector that enters our simulation: low-priced center-based care, high-priced center-based care, paid home-based care, and other paid care. Each model considers child age, income percentile, region, and number of parents as conditioning variables.⁹

During simulation, for each sector, each household is assigned a persistent price wedge drawn once at initialization from the predicted conditional distribution. This wedge remains constant across years and employment-care, ensuring that a household facing low out-of-pocket costs in one scenario (e.g. due to regional price variation) also faces low costs in counterfactual scenarios. For the "other paid" category (e.g. nannies, irregular care), which does not correspond to a formal supply sector in our model, we compute a weighted mean base price from NSECE household reports and estimate price wedges relative to this base. This allows us to assign realistic costs for this category even though it is not part of the equilibrium market-clearing process.

The next step is to expand each parent unit from one observed row to three rows representing counterfactual employment states: the primary caregiver working zero hours, part-time, or full-time. For each counterfactual employment state, we calculate counterfactual earnings by combining the primary caregiver's imputed hourly wage with employment-state-specific average hours and weeks worked. These averages are computed from the ACS sample: non-workers have zero hours, part-time workers average around 20--25 hours per week, and full-time workers average 40+ hours per week. The secondary earner's (non-primary caregiver's) labor supply remains fixed at observed levels across all counterfactuals. Counterfactual adjusted gross income (AGI) is calculated for each employment state by adjusting baseline AGI by the change in earnings. For parent units where parents file separately (unmarried cohabitants), we track AGI for each parent's tax unit separately. For married couples filing jointly, both parents' earnings contribute to a single joint AGI.

These counterfactual earnings and AGI values are used in simulation to calculate employment-state-specific tax liabilities, which in turn determine net income for each choice in the discrete choice model. To do this, we leverage data from The Budget Lab's Tax-Simulator model, a microsimulation model of individual income and payroll taxes. Because the ACS data lacks key variables (i.e. income composition, expenses for deductions, etc.), we do not directly calculate taxes for each parent unit using the Tax-Simulator's calculator logic. Instead, we perform a statistical matching exercise that pairs ACS records with Tax-Simulator records to get effective tax rates (ETRs) under different employment states.

For each tax unit record in Tax-Simulator, we calculate realized individual income effective tax rate and a series of counterfactual effective tax rates that would apply if earnings changed to different earnings levels across the population-wide distribution of wages. This schedule of ETRs allows us to approximate the tax schedule faced by the tax unit when earnings change, all else equal, which is precisely the experiment we run in the equilibrium simulator when we calculate net income under different employment choices (and thus different earnings).

The next step is to attach this information to the ACS data. Each ACS tax unit is matched to a donor from the Tax-Simulator pool based on observable tax-relevant characteristics. The matching procedure constructs cells of tax units from which to randomly sample, balancing precision against sparsity:

- Below the 90th AGI percentile, we match on AGI percentile bin, marital status, and number of children under age 13.
- Between the 90th and 98th AGI percentiles, we match on AGI percentile bin, filing status, and a binary indicator for whether the filer has any children under age 13.
- Above the 98th AGI percentile, we match on AGI percentile bin and filing status only.

Within each matching cell, donors are sampled with probability proportional to their survey weights. The matching is performed at the tax unit level (not the choice level) so that each tax unit receives a single donor assignment that applies across all employment choices. As described above, each donor tax unit comes with a schedule of effective tax rates at different wage levels corresponding to percentile thresholds ($p_0, p_5, \dots, p_{90}, p_{91}, \dots, p_{99}, p_{99.9}, p_{100}$). When calculating income tax liability for a given employment choice and thus wage, we linearly interpolate between the two nearest wage levels.

Payroll taxes (employee-side only) are calculated directly from counterfactual earnings using statutory rates. We apply the Medicare tax rate to all earnings and the Old-Age, Survivors, and Disability Insurance (OASDI) tax rate to earnings up to the OASDI wage cap for the simulation year.

Finally, taxes are allocated from the tax unit level (which defines tax liability) to the parent unit level (where choices are made in our model). For most parent units, these concepts are contiguous. For parent units with unmarried cohabitating parents, the parent unit's tax liability is the sum of each parent's tax liability.

NSECE Supply-Side Processing

In addition to surveying households about their ECEC use, the NSECE also conducts surveys of ECEC providers and workers. We draw on these surveys to estimate parameters for the supply side of our model: sector-specific prices, labor input ratios, and residual revenues.

We start with the center survey, which provides center- and classroom-level information on licensed centers, Head Start programs, and school-based pre-K. We take each center's reported weekly rate by age, binning into ages 0–1, 2–3, and 4–5, and calculating the implied hourly price by dividing by 40 hours. We also create a notion of full-time equivalent (FTE) enrollment by weighting part-time children to be one-half of full-time children. We compute the FTE enrollment-weighted median price within bins (excluding hourly rates of zero dollars). Then, we classify each center's age-specific care offering into one of three sectors as per the NSECE household survey classification described above: unpaid center-based care (zero price), high-priced center-based care (above-median price), low-priced center-based care (below-median price). We then aggregate to the sector level by taking the FTE enrollment-weighted mean by sector classification. The result of this process is sector-specific prices and enrollment totals.

The next survey we use is the workforce survey, which contains information on individual ECEC workers of center-based providers (and can be linked to specific providers in the center-based file). This survey gives us the information needed to calculate average hourly wages and effective labor input by sector and education. Workers are assigned a sector by linking to the center and classroom they work in. We calculate an FTE concept for workers by assigning a weight of 1 to full-time workers and 0.5 to part-time workers, then calculate the weighted FTE total by sector and education ("college" being defined as having at least a four-year degree and "non-college" as everyone else).

Combining these pieces, we calculate education-specific labor requirements (i.e. the reciprocal of hours-adjusted staffing ratios): for each sector, we divide college FTE employment by FTE enrollment and noncollege FTE employment by FTE enrollment. Per-hour labor costs by sector and education are calculated using average hourly wages—across *all* workers—from the worker sample. Residual revenue per hour is calculated as hourly price minus college hourly labor cost minus noncollege hourly labor cost.

We perform a similar series of calculations with a few modifications to get supply parameters for paid home-based care. We calculate the average hourly price for paid home-based care using the household survey by taking children observed using paid home-based care with valid net hourly costs and calculating a child-weighted mean net hourly price. We then use the home-based provider survey to construct staffing per enrolled child, where total workers is equal to one (the respondent/primary provider) plus number of reported helpers. Within the total staff count, we split workers into college and noncollege status based on the provider's self-reported education level and the response to a question about provider's reported percent of helpers with a four-year degree. Aggregating across providers yields total enrollment and total BA and non-BA workers, which we convert into labor requirements. Home-based labor costs are computed by applying the education-specific wage rates estimated from the center-based workforce survey to the home-based labor intensities. Residual revenue per hour is household-reported price minus this implied labor cost.

Across all calculations, we impose filters to ensure that sector totals are constructed from observations with internally consistent denominators and numerators. On the workforce side, we restrict to the main instructional roles (teachers/assistants) and require non-missing education and full-time/part-time status so that college/noncollege employment and FTE adjustments are well-defined; we also require a valid link to the sampled classroom age bin and to the center's sector classification so each worker is assigned to exactly one sector cell used in aggregation. On the center pricing/enrollment side, we exclude observations with missing or non-interpretable rates when computing the median and sector means. For home-based providers, we restrict to listed providers and we drop providers with missing enrollment/staffing/education information so staffing ratios are defined. Finally, for household-derived home-based prices, we restrict to children with valid net hourly cost measures.

Projections

The model uses 2019 as its base year because this is the most recent NSECE survey wave as of the time of writing. To project the microdata forward, we apply year-specific growth factors derived from historical data (for years before 2025) and CBO projections on macroeconomic and demographic interest. For each simulation year, we adjust household weights to reflect population growth by age and marital status. For monetary variables, we scale by relevant per-unit growth rates: hourly wage growth for hourly wages, GDP per capita for AGI.

Calibration

In this section, we discuss the calibration of our demand-side parameters using our 2019 ACS data (with ECEC choices imputed from the 2019 NSECE).

Utility Specification and Baseline Probabilities

Let N be the number of parent units in our data (with each observation having weight ω_i). Suppressing the time subscript t (and ignoring the utility shifter Δ_t , since we do not need to benchmark employment in the calibration year), recall that the deterministic component of parent unit i 's utility is given by:

$$V_{ij} = \alpha_{ij} + \beta \cdot \left(\frac{(Y_{ij} - T_{ij} - C_{ij})^{1-\rho}}{1-\rho} \right)$$

Here, again:

- Y_{ij} is i 's pretax income (AGI) associated with option j ;
- T_{ij} is i 's total tax liability associated with option j ;
- C_{ij} is i 's total cost of childcare associated with option j ;
- $\alpha_{ij} \in \mathbb{R}$ is an individual/option-specific intercept term; and
- $\beta \in \mathbb{R}_+$ and $\rho \in [0, 1)$ are utility parameters common to all parent units.

Let $\ell_{ij} = \ell_j \in \{\ell_{NW}, \ell_{PT}, \ell_{FT}\}$ be the (common-to-all- i) labor hours associated with each option j (with $\ell_{NW} = 0$). We take baseline choice probabilities P_{ij}^0 for each parent unit as exogenously given by our quantile random forest (QRF) model.

Calibration Overview and Elasticity Definitions

We need to calibrate a total of $N \times (J - 1) + 2$ parameters:

- The $N \times J$ intercept terms $\{\{\alpha_{ij}\}_{j \neq 1}\}_i$, with α_{i1} normalized to 0 for all i ; and
- The utility function parameters β and ρ .

We will make use of two targets from the literature:

- The **elasticity of maternal employment with respect to childcare costs** ε_C^E : the percent change in the aggregate maternal employment rate associated with a one-percent change in the price of paid childcare; and
- The **income elasticity of labor hours** ε_{ATT}^L : the elasticity of aggregate labor hours with respect to after-tax income.

For the former, following Brown et al. (2026) we use a value of -0.15 from [Morrissey \(2016\)](#). For the latter, we use a value of -0.05 following [CBO \(2012\)](#).

These two elasticity targets (for each model iteration), along with the $N \times J$ baseline choice probabilities $\{\{P_{ij}^0\}_j\}_i$ allow us to identify all of our parameters.

Backing Out Intercept Terms

In this section, we show that, given (β, ρ) and data, identification of the intercept terms can be accomplished with basic matrix operations (rather than, e.g., maximum likelihood estimation).

Define

$$\tilde{V}_{ij} \equiv V_{ij} - \alpha_{ij},$$

so that we have

$$P_{ij}^0 = \frac{\exp(\alpha_{ij} + \tilde{V}_{ij})}{\sum_m \exp(\alpha_{im} + \tilde{V}_{im})}$$

Take logs of both sides:

$$\ln P_{ij}^0 = \alpha_{ij} + \tilde{V}_{ij} - \ln \Sigma \iff \alpha_{ij} + \tilde{V}_{ij} - \ln P_{ij}^0 = \ln \Sigma$$

where $\Sigma \equiv \sum_m \exp(\alpha_{im} + \tilde{V}_{im})$. Thus for any $j, k \in \{1, \dots, J\}$:

$$\alpha_{ij} + \tilde{V}_{ij} - \ln P_{ij}^0 = \alpha_{ik} + \tilde{V}_{ik} - \ln P_{ik}^0 \iff \alpha_{ij} = \alpha_{ik} + (\tilde{V}_{ik} - \ln P_{ik}^0) - (\tilde{V}_{ij} - \ln P_{ij}^0)$$

Normalize $\alpha_{i1} = 0$.¹⁰ Then for any $j \in \{1, \dots, J\}$ we can write:

$$P_{ij} = \frac{\exp(\tilde{\alpha}_{ij} + \tilde{V}_{ij})}{\sum_m \exp(\tilde{\alpha}_{im} + \tilde{V}_{im})} = \frac{\exp(\tilde{\alpha}_{ij} + \tilde{V}_{ij})}{\sum_m \exp(\tilde{\alpha}_{im} + \tilde{V}_{im})} \cdot \frac{\exp(-\tilde{\alpha}_{i1})}{\exp(-\tilde{\alpha}_{i1})} = \frac{\exp(\tilde{\alpha}_{ij} - \tilde{\alpha}_{i1} + \tilde{V}_{ij})}{\sum_m \exp(\tilde{\alpha}_{im} - \tilde{\alpha}_{i1} + \tilde{V}_{im})} = \frac{\exp(\alpha_{ij} + \tilde{V}_{ij})}{\sum_m \exp(\alpha_{im} + \tilde{V}_{im})}$$

where $\alpha_{ij} \equiv \tilde{\alpha}_{ij} - \tilde{\alpha}_{i1}$ (and $\alpha_{i1} = 0$.) Then for any $j \neq 1$:

$$\alpha_{ij} = \alpha_{i1} + (\tilde{V}_{i1} - \ln P_{i1}^0) - (\tilde{V}_{ij} - \ln P_{ij}^0) = (\ln P_{i1}^0 - \ln P_{ij}^0) - (\tilde{V}_{ij} - \tilde{V}_{i1})$$

which depends only on (β, ρ) and our data.

Calibration Procedure

In what follows, given matrices \mathbf{A} we will write (with abuse of notation) $\ln(\mathbf{A})$ and $\exp(\mathbf{A})$, for the matrices formed by applying the scalar functions $\ln(\cdot)$ and $\exp(\cdot)$, respectively, to \mathbf{A} elementwise. Writing \mathbf{A}^{ex} indicates elementwise exponentiation (i.e. raising each element of \mathbf{A} to the power x). Finally, given a second matrix \mathbf{B} we write $\mathbf{A} \odot \mathbf{B}$ for the matrix formed via elementwise multiplication of \mathbf{A} and \mathbf{B} .

Initialization

1. Choose values for elasticity targets $(\varepsilon_C^E, \varepsilon_{ATI}^L)$ and a distance tolerance $\psi > 0$ (with ψ close to 0).
2. Form the following matrices (all $N \times J$):
 - The after-tax income matrix $\mathbf{ATI} \equiv ((Y_1 - T_1), \dots, (Y_N - T_N))'$, where $Y_i \equiv (Y_{i1}, \dots, Y_{iJ})$ and $T_i \equiv (T_{i1}, \dots, T_{iJ})$;

- The baseline probability matrix $\mathbf{P} \equiv (P_1, \dots, P_N)'$, where $P_i \equiv (P_{i1}^0, \dots, P_{iJ}^0)$;
- The normalized log-probability matrix

$$\mathbf{P}_{-1}^{\ln} \equiv \ln(\mathbf{P}) (\mathbf{I}_{J \times J} - \mathbf{Z}_{J \times J})$$

where $\mathbf{I}_{J \times J}$ is the $J \times J$ identity matrix and $\mathbf{Z}_{J \times J}$ is a $J \times J$ matrix with ones in the first row and zeroes in all other rows;¹¹

- The labor hours matrix $\mathbf{L} \equiv (L_1, \dots, L_i, \dots, L_N)'$, where $L_i \equiv (\ell_{i1}, \dots, \ell_{iJ})$;
- The employment matrix $\mathbf{E} \equiv (E_1, \dots, E_i, \dots, E_N)'$, where $E_i \equiv (\mathbf{1}_{\ell_{i1} > 0}, \dots, \mathbf{1}_{\ell_{iJ} > 0})$;
- The childcare cost matrix $\mathbf{C} \equiv (C_1, \dots, C_i, \dots, C_N)'$, where $C_i \equiv (C_{i1}, \dots, C_{iJ})$; and
- The paid childcare indicator matrix $\mathbf{D} \equiv (D_1, \dots, D_i, \dots, D_N)'$, where $D_i \equiv (\mathbf{1}_{C_{i1} > 0}, \dots, \mathbf{1}_{C_{iJ} > 0})$.

3. Form the $N \times 1$ vector of weights $\omega = (\omega_1, \dots, \omega_N)'$.

4. Choose a perturbation size $\delta > 0$ (with δ close to 0) and form the following perturbed matrices (again, all $N \times J$):

- The perturbed after-tax income matrix $\mathbf{ATI}^{\delta ATI}$, where

$$\mathbf{ATI}^{\delta ATI} \equiv \mathbf{ATI} + \delta \cdot [(\mathbf{P} \odot \mathbf{ATI}) \mathbf{1}_{J \times 1}] \mathbf{1}_{1 \times J}$$

where $\mathbf{1}_{n \times m}$ is an $n \times m$ matrix of ones;¹² and

- The perturbed childcare cost matrix $\mathbf{C}^\delta \equiv (1 + \delta) \cdot \mathbf{C}$.

5. Choose initial parameter values $(\beta^{(0)}, \rho^{(0)})$.

Iteration

For each $(\beta^{(t)}, \rho^{(t)})$:

1. **Compute intercepts.** Form the $N \times J$ matrix:

$$\tilde{\mathbf{V}}^{(t)} = \frac{\beta^{(t)}}{1 - \rho^{(t)}} \cdot \left((\mathbf{ATI} - \mathbf{C})^{\circ(1 - \rho^{(t)})} - \mathbf{1}_{N \times J} \right)$$

Then, compute the intercept matrix:¹³

$$\alpha^{(t)} = \mathbf{P}_{-1}^{\ln} - \tilde{\mathbf{V}}^{(t)} (\mathbf{I}_{J \times J} - \mathbf{Z}_{J \times J})$$

2. **Compute perturbed probabilities.** Using the intercept matrix $\alpha^{(t)}$ computed in the previous step, form the deterministic utility matrices associated with each of the perturbed matrices:

$$\mathbf{V}^{\delta ATI(t)} = \alpha^{(t)} + \frac{\beta^{(t)}}{1 - \rho^{(t)}} \cdot \left((\mathbf{ATI}^{\delta ATI} - \mathbf{C})^{\circ(1 - \rho^{(t)})} - \mathbf{1}_{N \times J} \right)$$

$$\mathbf{V}^{\delta C(t)} = \alpha^{(t)} + \frac{\beta^{(t)}}{1 - \rho^{(t)}} \cdot \left((\mathbf{ATI} - \mathbf{C}^\delta)^{\circ(1 - \rho^{(t)})} - \mathbf{1}_{N \times J} \right)$$

Then, form the corresponding perturbed probability matrices:¹⁴

$$\mathbf{P}^{\delta ATI(t)} = \text{diag} \left[\exp \left(\mathbf{V}^{\delta ATI(t)} \right) \mathbf{1}_{J \times 1} \right]^{-1} \exp \left(\mathbf{V}^{\delta ATI(t)} \right)$$

$$\mathbf{P}^{\delta C(t)} = \text{diag} \left[\exp \left(\mathbf{V}^{\delta C(t)} \right) \mathbf{1}_{J \times 1} \right]^{-1} \exp \left(\mathbf{V}^{\delta C(t)} \right)$$

3. Compute the income elasticity of labor hours. Compute the percent change in (weighted) aggregate labor hours caused by the ATI perturbation:¹⁵

$$\% \Delta L^{\delta ATI(t)} = \frac{\omega' \left[\left(\mathbf{P}^{\delta ATI(t)} - \mathbf{P} \right) \odot \mathbf{L} \right] \mathbf{1}_{J \times 1}}{\omega' \left[\mathbf{P} \odot \mathbf{L} \right] \mathbf{1}_{J \times 1}},$$

and compute the implied income elasticity of hours:

$$\varepsilon_{ATI}^{L(t)} = \frac{\% \Delta L^{\delta ATI(t)}}{\delta}$$

4. Compute the elasticity of maternal employment with respect to childcare costs. Compute the percent change in (weighted) aggregate employment caused by the perturbation:

$$\% \Delta E^{\delta C(t)} = \frac{\omega' \left[\left(\mathbf{P}^{\delta C(t)} - \mathbf{P} \right) \odot \mathbf{E} \right] \mathbf{1}_{J \times 1}}{\omega' \left[\mathbf{P} \odot \mathbf{E} \right] \mathbf{1}_{J \times 1}},$$

and compute the elasticity:

$$\varepsilon_C^{E(t)} = \frac{\% \Delta E^{\delta C(t)}}{\delta}$$

5. Compute the distance to targets and iterate. Calculate the Euclidean distance between the elasticities computed in the previous steps and the targets:

$$d^{(t)} = \sqrt{\left(\varepsilon_C^E - \varepsilon_C^{E(t)} \right)^2 + \left(\varepsilon_{ATI}^{L(t)} - \varepsilon_{ATI}^L \right)^2}$$

If $d^{(t)} > \psi$, update to $(\beta^{(t+1)}, \rho^{(t+1)})$ and repeat. Otherwise, set $(\beta, \rho) = (\beta^{(t)}, \rho^{(t)})$ and $\alpha_{ij} = \alpha_{i,j}^{(t)}$ for all i, j and stop.

Policy Reform Simulations

The model evaluates policy reforms by comparing equilibrium outcomes under a counterfactual policy scenario to a baseline scenario. This section describes the policy instruments available for simulation and the stochastic simulation procedure used to generate discrete household choices for a policy scenario after new equilibrium has been found.

Parameterization

Policy reforms enter the model through up to four distinct policy functions representing:

- **New demand-side subsidies.** These policies change families' out-of-pocket childcare costs. The policy function receives parent unit characteristics, market prices, and care costs, and returns a matrix of subsidy amounts for each parent unit-choice pair. This structure allows arbitrary subsidy designs indexed to any

observable household or choice characteristic. Examples include "copay" caps as a share of income or fixed-dollar grants that can be used to purchase certain kinds of care.

- **Supply-side policies.** These policies subsidize or mandate the production of ECEC. Supply-side policies can take the form of provider price subsidies, which create a government-funded wedge between consumer prices and provider prices (i.e., a subsidy of $\$S$ per care-hour means providers receive $\$(P + S)$ while families pay $\$P$; wage subsidies, where the government "tops up" ECEC worker wages by some per-hour rate based on credential level; or wage floors, which impose minimum ECEC worker wages by credential level and represent an unfunded mandate on providers. In each of these cases, these policies either pass through to lower market prices, higher wages, or some combination; the model does not assume a particular incidence pattern *ex ante*.
- **CDCTC policies.** The model includes a Child and Dependent Care Tax Credit (CDCTC) module that applies specified credit parameters to parent units (via tax units). This policy function receives tax unit characteristics (including income tax liability for accurate calculation under nonrefundable versions of the policy), market prices, and care costs *after* new demand-side subsidies, and returns a matrix of CDCTC amounts for each parent unit-choice pair.
- **Arbitrary tax credits.** These policies modify household resources independently of childcare decisions. The policy function adjusts tax liabilities (which can become negative for refundable transfers) for each employment state. Examples include child allowances, work credits with phase-ins, etc.

These policy functions are applied to each choice for each parent unit for each iteration during equilibrium calculation such that parent unit budget constraints and provider revenues reflect all policy changes.

Discrete Choice Draws

The equilibrium solver uses smooth softmax probabilities to aggregate demand. After solving for equilibrium prices, the model simulates discrete household choices by sampling from each household's distribution of random utility. This discretization is required for our estimates of impacts on children's future earnings, which looks at gross flows between ECEC arrangements (rather than expected values).

For each parent unit, independent Gumbel(0, 1) shocks (i.e. Type I extreme value) are drawn for every choice in the choice set. These shocks are held constant across scenarios. For each parent unit record, the model computes total utility as the sum of the deterministic component and the stochastic shock; the household's discrete choice is the alternative with highest total utility.

Model Output

Upon solving the new equilibrium under a policy reform, the model produces a number of metrics summarizing the impact of the reform. These metrics highlight the economic, fiscal, and budgetary tradeoffs associated with ECEC policy design.

Employment and ECEC Decisions

The model tabulates the distribution of children across ECEC arrangements, grouped by child age and care intensity (i.e. part-time or full-time), under both the baseline and the policy scenario. The model also reports parental employment counts and rates by sex and intensity (part-time or full-time).

Fiscal

The model calculates annual budgetary effects by comparing policy and baseline scenarios. The total budget effect aggregates three channels for budgetary impact: changes in demand subsidies (including both household and provider subsidies), CDCTC changes (e.g. changes in credit eligibility due to ECEC arrangement changes), and other individual income and payroll tax changes (e.g. changes in tax liability due to changes in employment).

We decompose fiscal effects into mechanical and behavioral components. The mechanical effect represents the policy's cost if household behavior remained exactly at baseline—that is, applying the policy's subsidy rules to baseline employment and care choices, evaluated at baseline market prices. This isolates the direct fiscal impact from policy generosity. Subtracting the mechanical cost from the total cost yields the behavioral component, reflecting the portion of the cost attributable to behavioral responses (e.g. ECEC type switching and parental employment changes) and market price adjustments. Mechanical effects are useful for understanding the first-order generosity and targeting of a policy reform, while total effects capture the policy's actual estimated impact.

Distribution

The model calculates distributional impacts by comparing family-level outcomes under policy and baseline scenarios across multiple dimensions: percentiles and levels of AGI, parent age, marital status, and number of children under age five.

The core distributional metric is change in net income, defined as the change in AGI less taxes (net of CDCTC) and out-of-pocket ECEC costs. As with the fiscal analysis, distributional impacts are decomposed into mechanical and behavioral components. This is especially important for distribution analysis, where the goal is typically to measure a notion of change in *welfare*, not change in income after behavioral responses, especially in settings where a policy causes prices to change.

For example, a family ineligible for a new subsidy who unenrolls from daycare as a result of higher prices will see an increase in net income after the switch. But they are now worse off, since they valued that daycare option at more than the old price but can no longer afford it. To that end, we produce three different concepts for each distribution metric, each designed to measure something different:

- **Mechanical effect.** This is change in net income when the policy is applied at baseline market prices and families maintain their baseline employment and ECEC choices. It captures the "first-order" or "static" changes associated with a policy.
- **Welfare effect at policy prices ("economic value").** This is a second-order approximation to the Hicksian equivalent variation, ignoring income effects (which are calibrated to be small in our model). This measure is computed as the average of two terms: (i) the change in net income when the policy is applied at new equilibrium prices but families maintain their baseline choices, and (ii) the change in net income when baseline rules are applied at baseline prices but families are evaluated at their policy choices. Formally, if k_0 and k_1 denote baseline and policy choices, p_0 and p_1 denote baseline and policy rules, and P_0 and P_1 denote baseline and equilibrium prices:

$$\text{Welfare} = \frac{1}{2} [\text{NI}(p_1, P_1, k_0) - \text{NI}(p_0, P_0, k_0)] + \frac{1}{2} [\text{NI}(p_1, P_1, k_1) - \text{NI}(p_0, P_0, k_1)]$$

This is our preferred welfare measure because it captures how much better or worse off families are given the new economic environment they face, without being sensitive to the particular reference bundle at which the price and policy changes are evaluated.

- **Total effect.** This is the final change in net income after all behavioral and price changes. This metric, while not a measure of welfare, is useful for looking at the distribution of ECEC cost burdens before and after a policy is enacted.

For each concept, we produce several metrics: change in net income, percent change in after-tax income, share of families who are better off, share of families who are worse off, and ECEC burden as a share of AGI.

Children's Future Earnings

Nonparental childcare participation in early childhood influences children's human capital development and affects outcomes later-in-life. In post-processing, our ECEC model converts policy-driven shifts in childcare allocations into changes in early test scores and subsequent later-in-life earnings.

Imputing early test scores

The change in a child's test score as a result of policy reform depends on the type of care that child receives in both baseline and policy. For example, we might expect that children moving from parental care to high-cost center-based care see larger changes than those moving from unpaid center-based care to high-cost center-based care. Imputing test score effects thus requires childcare for each pair of baseline and scenario childcare options.

The empirical literature, however, almost entirely reports the test score gains from center-based care against a mix of counterfactual choices. From our review of the literature, we ultimately include sixteen estimates of the effects on test scores from across eight studies in our imputations process (see the Appendix for a discussion of included studies). At a high-level, we take the following steps to back out the pairwise treatment effects implied by each study:

1. **Solve for implied effects by counterfactual care choice.** For each estimate, we decompose the effect of center-based care into separate effects for children drawn from home-based, other (e.g., relative care), and parental care.
2. **Construct a matrix of treatment effects.** Aggregating our above results, we construct a 4×4 matrix of treatment effects where cell (i, j) contains the effect of switching from the i th childcare option to the j th.
3. **Account for quality differences by price.** Lastly, we adjust the effects for center-based, home-based, and other care options to reflect differences in quality by price group, resulting in a final 8×8 matrix of treatment effects.

We conduct this process separately for children ages 3–4 and children younger than age 3, making minor adjustments as needed. The remainder of this section describes our approach in more detail.

Step 1: Implied effects by counterfactual care choice. By the law of iterated expectations, the LATE associated with center-based care, which we will denote $LATE_c$, can be decomposed into a weighted average of sub-LATEs measuring the effect with respect to each counterfactual care option:

$$LATE_c = \sum_d s_{d,c} LATE_{d,c}, \quad \text{where } d \in \{h, o, p\},$$

where subscripts $\{c, h, o, p\}$ refer to center-based, home-based, other, and parental care, respectively, $s_{d,c}$ denotes the share of compliers who switch from one of home-based, other, or parental care into center-based care, and

$LATE_{d,c}$ denotes the corresponding LATE for that sub-group.

Most studies only provide estimates of $LATE_c$ and the complier shares. For studies of children younger than age 3, only two studies decompose treatment effects by the alternative childcare option. Taking the inverse-variance weighted (IVW) average of the $LATE_{h,c}$ estimates and the IVW average of the $LATE_{o,c}$ estimates from these two studies as given, we compute for the $LATE_{p,c}$ implied by each remaining study. We then use a parametric bootstrap to compute standard errors.¹⁶ Finally, we take an IVW average across the implied $LATE_{p,c}$ values.

Step 2: Treatment effect matrix. Step 1 results in an incomplete 4×4 treatment-effects matrix, where cell (i, j) contains the effect of switching from the i th childcare option to the j th:

$$\begin{bmatrix} LATE_{c,c} & LATE_{c,h} & LATE_{c,o} & LATE_{c,p} \\ LATE_{h,c} & LATE_{h,h} & LATE_{h,o} & LATE_{h,p} \\ LATE_{o,c} & LATE_{o,h} & LATE_{o,o} & LATE_{o,p} \\ LATE_{p,c} & LATE_{p,h} & LATE_{p,o} & LATE_{p,p} \end{bmatrix}.$$

The bold entries are calculated during our above process. To fill out the remaining entries, we assume that local average treatment effects are additive across different groups of compliers; i.e.,

$$LATE_{x,y} = LATE_{x,z} - LATE_{z,y}, \quad \text{for } x, y \in \{c, h, o, p\}.$$

This step effectively assumes that potential outcomes for different groups of compliers are sufficiently similar such that the right-hand sum of conditional expectations equals the conditional expectation of the sum. While this is a strong assumption, it yields a treatment-effects matrix with intuitive properties: notably, our matrix is skew symmetric (the effect of switching between the same type of care is zero and inverted moves are opposite in sign).

Step 3: Quality adjustments by price. Lastly, we disaggregate childcare options by price (high-priced, low-priced, and unpaid center-based care and paid/unpaid home-based and other care) for a final 8×8 treatment-effects matrix. We do this by adjusting the bottom row (i.e., switches from parental care) of the above 4×4 treatment-effects matrix for differences in childcare quality by price.

Using the 2019 NSECE center-based (and workforce) and home-based surveys, we construct a quality index as the simple average of the following five variables (normalized such that $\min(x) = 0$ and $\max(x) = 1$):

- Whether the respondent uses curriculum,
- Whether the respondent has a bachelor's degree,
- The classroom's staff-to-student ratios,
- One minus the provider's departure rate (home-based providers do not report this information), and
- Whether the provider offers meals.

As an illustrative example, suppose that the quality of high-priced, low-priced, and unpaid centers are given by q_{hc} , q_{lc} , and q_{uc} , respectively. We compute quality-adjusted sub-LATEs such that:

- Effects scale relative to quality:

$$\text{LATE}_{p,hc} = \frac{q_{hc}}{q_{uc}} (\text{LATE}_{p,uc}),$$

- The quality-adjusted sub-LATEs average to the unadjusted LATE from above:

$$\frac{1}{3} (\text{LATE}_{p,hc} + \text{LATE}_{p,lc} + \text{LATE}_{p,uc}) = \text{LATE}_{p,c}.$$

We repeat a similar adjustment for home-based care. This quality adjustment yields the bottom row of our final 8 × 8 treatment-effects matrix. Following the earlier assumption of additivity, we can again fill out the remainder of our treatment-effects matrix.

Differences for preschool-age children.

For children ages 3–4, studies estimate LATEs for a specific type of unpaid center (e.g. Head Start), and those reporting sub-LATEs only decompose their estimate by compliers from other centers and compliers from non-centers. Taking the former sub-LATE as given, we repeat a similar process as described above to compute the latter effect across studies. We then assume that the sub-LATE for compliers from non-centers accurately captures $\text{LATE}_{p,uc}$. We perform a similar quality adjustment to compute the effects for high- and low-priced centers. Since we lack information for home-based and other care for 3–4-year-olds, the remainder of the bottom row uses the same LATEs as in the matrix for children under 3 years old.

Projecting Later-in-life Earnings

Unlike the pairwise test score effects described above, the relationship between early test scores and later-in-life earnings is well documented. [Chetty et al. \(2011\)](#) find that a 1 standard deviation increase in test scores (from improvements in the early childhood care quality) increases average earnings at age 27 by 13.1 percent. As such, the percent change in later-in-life earnings for a child switching from care option x to option y is given by $0.131 \times \text{LATE}_{x,y}$. To compute the dollar value of this change, we compute average child baseline earnings at age 27 based on parent income rank (see [here](#) for more details).

Net Present Fiscal Value

We compute a net present value (NPV) measure comparing the lifetime fiscal returns from children’s increased future earnings against the fiscal cost of a policy reform. The idea is to estimate the extent to which subsidies for early childhood care “pay for themselves” through higher tax revenues later. We do so by following one cohort of children who were exposed to a full five years of a policy change starting in year t .

We approximate the cohort’s associated fiscal costs by computing the cross-sectional costs by child age as a proxy for a single cohort’s historical costs, adjusted for time. Under steady-state assumptions, this is equivalent to tracking one cohort longitudinally.

For each child age group $a \in \{0, 1, 2, 3, 4\}$ in year t , we compute the net direct fiscal cost C_a (inclusive of tax changes due to parental labor supply changes, expressed in real 2025 dollars) attributable to aged- a children (assuming equal-split of fiscal costs in two-child parent units). Costs are compounded forward to age 4 (the analysis cohort’s age) using the real discount rate r :

$$\text{Cost}(PV) = \sum_{a=0}^4 C_a (1+r)^{(4-a)}.$$

On the returns side of the question, we:

1. **Identify child-level earnings impacts as per the discussion above.** For each child, we compare ECEC arrangements under baseline and policy scenarios. A child is a “switcher” if their care type changes, and we apply the associated earnings effect from the LATE matrix, yielding a percent change in earnings g . We assume this change represents a relative level shift in earnings at all ages relative to baseline.
2. **Restrict to age-4 children.** Returns are computed only for children at age 4, the final year of the model’s ECEC window. This assumption, which we view as a conservative approach, reflects several factors:
 1. The model is fundamentally cross-sectional; we cannot track each child’s path of ECEC arrangements. If we were to sum across all child-year observations of switchers, effects for some children would be double counted. For example, some simulated 4-year-olds who are induced to switch into a specific ECEC type in year t may also have switched as 3-year-olds in year $t - 1$. If test score impacts are cumulative and not annual (i.e., treatment is binary rather than “dose-response”), then adding effects across each child-year would represent double counting.
 2. Test score gains from interventions at younger ages tend to fade out, so focusing on age 4 captures the persistent effects we care about.
 3. Children treated at younger ages are likelier to remain in their care arrangement through age 4, so this cohort implicitly reflects cumulative treatment.
3. **Calculate mean earnings impacts for each parent income rank.** For each parent income percentile $j = 1, 2, \dots, 100$, we compute the mean earnings percent change g_j among age-4 switchers at that rank. We also track n_j , the total number of age-4 children at rank j .
4. **Calculate implied change in taxes.** For each future year k from $t + 27$ through $t + 65$, we read projected mean earnings $wL_j(k)$ and dollar-weighted mean marginal tax rates $\tau_j(k)$ from Tax-Simulator output for individuals at parent rank j in year k .¹⁷ The marginal tax rate includes both income and payroll (both employee- and employer-side) taxes. We calculate the tax change at rank j as:

$$\Delta T_{jk} = wL_{jk} g_j \tau_{jk} n_j.$$

5. **Discount and aggregate.** We discount these amounts back to year t using the discount rate r and inflation rate π :

$$\Delta T_{jk}(PV) = \frac{\Delta T_{jk}}{(1 + r + \pi)^{(k-t-4)}}.$$

Then sum over all values of parent income rank j and year k .

These processes give us both net direct cost and net fiscal return for a given cohort, both in present value terms. We calculate a “net fiscal recovery ratio” equal to the ratio of PV returns to PV costs.

The model computes present value metrics using discount rates ranging from 1 to 10 percent; results at the 3 percent rate (the reference rate in [Hendren and Sprung-Keyser \[2020\]](#)) are reported as the primary estimate. Projected inflation rates are year-specific and come from CBO’s Economic Outlook.

Technical Implementation

The model is implemented in R as a modular microsimulation pipeline organized into four phases: (1) input processing, which transforms raw NSECE and ACS microdata into the hierarchical parent-unit data structures described above and trains the imputation models; (2) calibration, which estimates the demand model's preference parameters to match target elasticities; (3) simulation, which solves for equilibrium prices under baseline and counterfactual policy scenarios using L-BFGS-B optimization, then draws discrete household choices via simulated Gumbel shocks; and (4) output aggregation, which computes the fiscal, distributional, and child-outcome metrics described in the preceding sections. Policy reforms are specified via pluggable R functions loaded at runtime from configuration directories, allowing new subsidy designs to be evaluated without modifying the codebase.

The simulator is driven by CSV "runscripts" that define scenario configurations—each row specifies a demand-side policy, a supply-side policy, a demand model, a supply model, a wage floor, and the years to simulate enabling comparisons of multiple policy designs in a single execution. Because the equilibrium-finding process can be computationally intensive at larger ACS sample, the model offers parallelization across years via a SLURM cluster.

The codebase is available on the [Budget Lab's GitHub page](#).

Limitations and Next Steps

Our results should be interpreted in light of several modeling choices driven by computational tractability and data constraints. This section discusses the assumptions we view as most consequential for understanding our results and outlines directions for future work.

- **Limited measure of income.** The model uses adjusted gross income (AGI) as its principal measure of family resources. AGI excludes transfer income (SNAP, SSI, housing vouchers), non-market sources of financing consumption like familial transfers, and other non-taxable sources. For low-income families, this understates available resources and thus overstates the marginal utility gain from entering employment, potentially inflating modeled employment responses to ECEC subsidies. Beyond mismeasurement of income and consumption, the absence of fiscal transfers presents two other issues. First, transfer phase-outs add to effective marginal tax rates for low-income families, meaning the model understates the implicit cost of entering employment for these families. Second, on the fiscal side, increased parental employment reduces transfer outlays—a budgetary offset the model does not count. The calibration process partially ameliorates this issue: because the model's preference parameters are fit to match aggregate employment elasticities, the calibrated parameters will implicitly absorb the "missing" disincentive from transfer phase-outs by producing a weaker baseline taste for employment than the true structural preference. But if new ECEC policies interact with these effects (that is, the effects cannot be "differenced out" across scenarios), then absorption into the preference terms will bias behavioral estimates.
- **Incomplete modeling of baseline subsidies.** The model does not explicitly represent existing ECEC subsidy programs, for example CCDF vouchers. Instead, baseline subsidy receipt is absorbed into the household-specific price wedge estimated from NSECE data. This means the model cannot capture the fiscal interaction between existing and new subsidies: if a policy reform raises market prices, the cost of existing income-tested subsidy programs may also increase. We plan to explore the viability of isolating the component of the price wedge attributable to regional variation or quality differences, then modeling current-law subsidies explicitly.

- **Simplified supply technology.** The supply side makes several assumptions that, taken together, limit the model's ability to capture the full range of provider responses to policy.

First, production uses a fixed-proportions technology, where each hour of care requires a fixed ratio of college-educated and non-college-educated workers. Providers cannot substitute between worker types, so a policy targeting one credential group (e.g. a wage subsidy for credentialed workers) passes through to prices at exactly the calibrated labor intensity, with no margin for staffing adjustment.

Second, the supply-side equilibrium condition pins per-unit residual revenue—the gap between the price received by providers and their unit labor cost—to an assumption value consistent with observed data in 2019. This is analogous to a zero-economic-profit condition with a fixed markup, and it rules out the possibility that providers earn higher or lower margins when demand shifts. All adjustment to demand shocks thus flows through wages and prices rather than profits. Whether this assumption is warranted likely depends on the policy in question: it is clearly more appropriate for policies that explicitly include cost-control mechanisms (such as setting provider reimbursement rates based on a cost structure formula rather than reported sticker prices for care, as has been the case in several legislative proposals for federal childcare subsidies) than for policies that naively reimburse whatever sticker price that providers charge.

Third, the model relies on the assumption that flexible sticker prices are the mechanism by which the ECEC market re-equilibrates after a policy change. In practice, capacity constraints, waitlists, and lottery-based admission are common features of ECEC markets, particularly for publicly funded programs like Head Start and state pre-K. These policies, which lack a well-defined notion of a "sticker price" that can rise to meet demand, rely on quantity rationing when demand exceeds available slots. We plan to extend the model to allow for rationing as a form of market clearing.

- **Abstraction away from multiple care arrangements.** Our discrete choice framework limits each child to one care arrangement. In reality, parents often place their children in multiple care arrangements (for example, half-day preschool in the morning and informal care provided by a grandparent in the afternoon). As described in Section 3.1, for children with multiple reported care arrangements, we apply an aggregation rule that limits each child to one arrangement. Put differently, we model what is a continuous allocation decision (which admits a mix of extensive and intensive margin responses to policy changes) as a discrete allocation decision (where all care choices are extensive-margin only). This assumption loses information about baseline enrollment patterns and may bias the estimates of policy changes which change the relative attractiveness of choosing multiple care arrangements and/or operate principally upon intensive care usage margins.
- **Missing human capital channel for parents.** The model captures the contemporaneous earnings and fiscal effects of changes in maternal employment but treats wages as fixed across employment states. In practice, mothers who maintain labor force attachment accumulate experience, avoid skill depreciation, and may follow steeper earnings trajectories—that is, ECEC policy has the potential to reduce maternal wage caps. By ignoring this channel, the model could understate the long-run earnings and fiscal benefits of policies that increase parental employment.

In our review of the literature, we prioritized studies that leverage a research design based on random assignment and estimate local average treatment effects (LATEs), as these studies typically produce the highest-quality evidence of causal effects.¹⁸ We limit our review to studies of more recent birth cohorts (within the last three decades) in the United States as improvements in counterfactual childcare options over time make evidence from older cohorts less relevant for policy evaluations today ([Duncan and Magnuson, 2013](#); [Cascio, 2021](#)).¹⁹ Estimates from our eight included studies are summarized below.

Effect of Childcare Participation on Early Test Scores

Study	Design	Treatment	Counterfactual Care	Age	Cohorts	Est.
Kline and Walters (2016)	Lottery	Head Start	Non-center	3-4 yo	2002-2004, entry	0.37
Kline and Walters (2016)	Lottery	Head Start	Other center	3-4 yo	2002-2004, entry	-0.093
Feller et al. (2016)	Lottery	Head Start	Non-center	3-4 yo	2002-2004, entry	0.23
Feller et al. (2016)	Lottery	Head Start	Other center	3-4 yo	2002-2004, entry	0
Cascio (2023)	DD/IV	UPK ¹	Mix	3-4 yo	2001, birth	0.569
Cascio (2023)	DD/IV	TPK ¹	Mix	3-4 yo	2001, birth	-0.373
Lipsey et al. (2018) ³	Lottery	TPK ² (TN)	Mix	3-4 yo	2009-2011, entry	0.266
Humphries et al. (2025)	Lottery	UPK ¹ (New Haven, CT)	Mix	3-4 yo	2003-2022	0.062
Herbst (2013)	IV	Center-based	Parental	<3 yo	2001, birth	-0.407
Herbst (2013)	IV	Center-based	Home-based	<3 yo	2001, birth	-0.483
Herbst (2013)	IV	Center-based	Other	<3 yo	2001, birth	-0.242
Bernal and Keane (2011)	IV	Center-based	Parental	<3 yo	1990-2000, birth	0.101
Bernal and Keane (2011)	IV	Center-based	Home-based	<3 yo	1990-2000, birth	0.147
Bernal and Keane (2011)	IV	Center-based	Other	<3 yo	1990-2000, birth	-0.161
Herbst (2025)	Lottery	Center-based	Mix	<3 yo	1990-1996, entry	0.3

¹ *Universal Pre-k*; ² *Targetted Pre-k*; ³ *Also in Watts et al. (2019)*

Created with [Datawrapper](#)

Meanwhile, quality evidence of the effects of childcare attendance on later-in-life earnings is limited. While older evidence from demonstration programs during the 1960s and 1970s documents sizable effects ([Elango et al., 2016](#)), these studies are less likely to be policy relevant today. Because of changes in counterfactual care options alongside differences in program quality between these small scale demonstration programs and current childcare providers, the magnitude of these estimates on later-in-life earnings would likely be unrealistic.

Predicting earnings from early test scores allows us to more accurately capture the effects of childcare in a more contemporary context and is common practice in the childcare literature. [Kline and Walters \(2016\)](#) and [Cascio \(2023\)](#) likewise leverage results on later-in-life earnings from [Chetty et al. \(2011\)](#) in cost-benefit analyses of Head Start and universal pre-K (UPK), respectively. Moreover, because test score gains only capture one pathway through which childcare might impact later-in-life outcomes, this approach likely understates the effects on earnings, effectively providing a lower bound.

Studies examining other outcomes suggest that approximations using early test score gains may underestimate later-in-life benefits. For reductions in crime, [Anders et al. \(2023\)](#) show that actual effects are more than 7 times greater than the implied outcomes from Kindergarten test scores. Performing this exercise for college enrollment similarly reveals underestimates. Back-of-the-napkin calculations using results from [Chetty et al. \(2011\)](#) imply that a 1 SD increase in test scores increases college attendance by 4.1pp. Using the reported early test score gains in [Kline and Walters \(2016\)](#) and [Cascio \(2023\)](#), imputed college enrollment effects yield an increase of 2.3pp/1.5pp for UPK/Head Start. Studies that directly estimate college enrollment gains, however, find increases of 8.3~pp for UPK participation ([Gray-Lobe et al., 2023](#)) and 5.4~pp for Head Start participation ([Bailey et al., 2021](#)). In both cases, implied effects are around 70% smaller.

Footnotes

- 1 Brown, Jessica H., Chloe R. Gibbs, Neale Mahoney, and T.V. Ninan. 2026. "The Stanford Childcare Model." Palo Alto, CA: Stanford Institute for Economic Policy Research.
- 2 We are grateful for advice from and discussions with the authors of those papers, which were invaluable in the development of The Budget Lab's model.
- 3 Nonrandom missingness is only a problem in our framework to the extent it cannot be explained by observable characteristics that enter our imputation model, which estimates ECEC choice behavior as a function of observable characteristics and then fits values on non-missing data in the ACS.
- 4 Because we are principally interested in simulating the effects of policies that subsidize paid ECEC arrangements, it is more appropriate to rank by this notion of priority and not by hours. For example, a child who spends two days a week in a home-based daycare but three days with a relative should, for our purposes of understanding the ECEC market, be classified as having the former as their "primary" arrangement.
- 5 Technically, the impact of the number of parents is reflected not via an X variable but rather via separately estimating the model by the number of parents. This is done so that sex of primary earner is given an internally coherent definition across single and joint parent units.

- 6 While considering all features at each split is often said to risk overfitting, in our case, the exercise is not a traditional out-of-sample forecasting exercise but rather a statistical matching exercise where we are taking the donor dataset (NSECE) as "ground truth" from which to sample donor records. In other words, "overfitting" is not a well-defined concept in our application.
- 7 Put differently, the ECEC imputation random forest returns $p(\text{ECEC choice} \mid \text{employment choice})$, which is then multiplied by $p(\text{employment choice})$ to generate $p(\text{ECEC choice} \cap \text{employment choice})$.
- 8 For example, the impact of a policy that caps ECEC care costs would be overestimated in a "mean-only" pricing assumption when the mean cost exceeds the cap but some families face costs below the cap.
- 9 The specific inclusion of income and family-structure variables in these models is partly intended to capture baseline price variation due to means-tested subsidies.
- 10 This normalization is without loss of generality: suppose the "true" intercept terms were $\{\tilde{\alpha}_{i1}, \dots, \tilde{\alpha}_{iJ}\}$. Then for any $j \in \{1, \dots, J\}$ we can write:

$$P_{ij} = \frac{\exp(\tilde{\alpha}_{ij} + \tilde{V}_{ij})}{\sum_m \exp(\tilde{\alpha}_{im} + \tilde{V}_{im})} = \frac{\exp(\tilde{\alpha}_{ij} + \tilde{V}_{ij})}{\sum_m \exp(\tilde{\alpha}_{im} + \tilde{V}_{im})} \cdot \frac{\exp(-\tilde{\alpha}_{i1})}{\exp(-\tilde{\alpha}_{i1})} = \frac{\exp(\tilde{\alpha}_{ij} - \tilde{\alpha}_{i1} + \tilde{V}_{ij})}{\sum_m \exp(\tilde{\alpha}_{im} - \tilde{\alpha}_{i1} + \tilde{V}_{im})} = \frac{\exp(\alpha_{ij} + \tilde{V}_{ij})}{\sum_m \exp(\alpha_{im} + \tilde{V}_{im})},$$

where $\alpha_{ij} \equiv \tilde{\alpha}_{ij} - \tilde{\alpha}_{i1}$ (and $\alpha_{i1} = 0$).

- 11 Multiplying an $n \times m$ matrix \mathbf{A} by $(\mathbf{I}_{m \times m} - \mathbf{Z}_{m \times m})$ is the same as subtracting the first column of \mathbf{A} from each column of \mathbf{A} .
- 12 This is matrix notation for adding δ times $\sum_j P_{ij}^0 ATI_{ij}$ to ATI_{ij} for all i, j - in other words, adding a small percentage of *expected* baseline ATI to every option for every parent unit.
- 13 This is the matrix version of the final equation of Section 4.2 above.
- 14 Multiplying an $n \times m$ matrix \mathbf{A} by $\text{diag}[\mathbf{A}\mathbf{1}_{m \times 1}]^{-1}$ is the same as taking each element of a given row of \mathbf{A} and dividing it by the sum of elements in its row.
- 15 For an $n \times m$ matrix \mathbf{A} and an $n \times 1$ vector of row weights ω , the operation $\omega' \mathbf{A}\mathbf{1}_{m \times 1}$ corresponds to computing a weighted sum of all of the elements of A , with the elements in row i weighted by ω_i . Here, for example, the denominator of the fraction is equal to $\sum_i \omega_i \left(\sum_j P_{ij}^0 \cdot \ell_{ij} \right)$ - or, in other words, baseline aggregate labor hours.
- 16 For each of the remaining studies, we simulate 10,000 draws from a normal distribution implied by the estimate for LATE_c and by the IVW averages of $\text{LATE}_{o,c}$ and $\text{LATE}_{h,c}$ (assuming independence across distributions). Then for complier shares, we use a multinomial distribution where the size parameter is chosen such that the average standard error across the implied marginal distributions matches the average of the reported standard errors for complier shares.
- 17 For more information on how we impute parent income rank in Tax-Simulator data, refer to [this report](#).
- 18 These results are also largely consistent with studies of universal pre-k programs in Oklahoma ([Gromley and Gayer, 2005](#)), Georgia ([Peisner-Feinberg et al., 2014](#)), and Boston ([Weiland and Yoshikawa, 2013](#)) that utilize age-cutoff regression discontinuity designs (see [Lipsey et al. \[2015\]](#) for a broader discussion of the results and limitations of this design).
- 19 Notably, this restriction excludes evidence from demonstration studies during the 1960s and 1970s (e.g. The Perry Preschool Project).