Review of Economic Dynamics 56 (2025) 101270

Contents lists available at ScienceDirect

Review of Economic Dynamics

journal homepage: www.elsevier.com/locate/red

Full Length Article

Family policies and child skill accumulation ☆

Emily G. Moschini [a,*], Monica Tran-Xuan [b]

[a] William & Mary, Department of Economics, 258 Chancellors Hall, Williamsburg, VA 23185, United States of America
[b] SUNY at Buffalo, Department of Economics, 425 Fronczak Hall, Buffalo, NY 14260, United States of America

ARTICLE INFO

JEL classification:
J13
J18
J24
D64

Keywords:
Childcare subsidy
Child tax credit
Early childhood
General equilibrium
Skill investment
Welfare

ABSTRACT

We analyze the economic effects of two major family policies in the United States, the Child Tax Credit and the Child Care and Development Fund childcare subsidy, in an overlapping generations framework where altruistic parents invest in their child's skill using their own time and purchased childcare time. The model incorporates differences in the design of these policies and endogenizes low rates of childcare subsidy receipt by including application costs and subsequent rationing. We compare the effects of a recent child tax credit expansion with a spending-equivalent expansion of the childcare subsidy implemented by reducing access frictions. Across steady states, the childcare subsidy expansion generates a larger increase in average adult skill, which leads to larger welfare gains behind the veil of ignorance compared to the tax credit expansion. However, the two policies yield similar average welfare gains for adults who know their own skill level, and the tax credit benefits a larger share of this group.

1. Introduction

Economists have long acknowledged that aid to families with children can bolster children's accumulation of skill and improve the child's labor market outcomes later in life (Leibowitz, 1974). Indeed, "family policies"—that is, government transfers whose receipt or intensity is tied to the presence of children—have been common across countries in recent decades (Kamerman, 2000; Olivetti and Petrongolo, 2017). In the United States, family policies may be implemented through the tax system or via a spending program. In the latter case, a bureau evaluates applicants and grants access to aid after determining eligibility; in practice, spending programs entail a frictional process of application and receipt, and few among the eligible receive aid (Sommartino et al., 2002; Johnson et al., 2011).

In this paper, we focus on two major US family policies: the Child Tax Credit (CTC/ACTC, hereafter CTC), a cash transfer which is targeted to middle-income families and implemented via the tax system, and the Child Care and Development Fund (CCDF), a childcare subsidy which is targeted to the poor and implemented via a spending program. In recent years, annual spending on the CTC is higher than spending on the CCDF: \$51.5 billion versus \$5.5 billion per year, or 0.27 percent versus 0.03 percent of output,

☆ We are grateful to the editor and three anonymous reviewers for their constructive comments. We also thank Orazio Attanasio, Adam Blandin, Katarina Borovickova, Mary Ann Bronson, Bettina Brüggemann, Michael Carter, Diego Daruich, Nezih Guner, Christopher Herrington, Zachary Mahone, Andrea Moro, Joseph Mullins, Sergio Ocampo Díaz, Gajendran Raveendranathan, Todd Schoellman, Nicolas Trachter, and Chao Wei for helpful comments, as well as seminar participants at McMaster University and conference participants at LA Macro (Greenville, SC), I-85 Workshop (Richmond, VA), National Tax Association (Denver, CO), Southern Economics Association (New Orleans, LA), Midwest Macro Association (Chicago, IL), and the Society for Economic Dynamics (Barcelona, Spain). Qiong Liu provided help with the Panel Study of Income Dynamics for an earlier version of this paper. All errors are ours.

* Corresponding author.
E-mail addresses: egmoschini@wm.edu (E.G. Moschini), monica.tranxuan@gmail.com (M. Tran-Xuan).

https://doi.org/10.1016/j.red.2025.101270
Received 8 August 2023; Received in revised form 26 December 2024

Available online 9 January 2025

respectively.[1] This difference in spending is due to differences in eligibility rules, transfer levels for recipients, and access frictions associated with the CCDF's implementation method such as application costs and waiting lists (Guzman, 2019). Starting from this policy baseline, we examine the tradeoffs of expanding one policy versus the other. The specifics of the CTC expansion are drawn from the 2017 Tax Cuts and Jobs Act (TCJA), while the CCDF expansion reduces the frictions associated with application and receipt for that program. The two expansions result in the same change in total government spending as a percent of output.

Specifically, we embed representations of the CTC and CCDF into an overlapping generations model. Early in adulthood, all consumers become parents who altruistically invest in their child's initial skill using their own time and purchased childcare time; the child's initial skill is a stochastic draw for the parent, and the model allows for intergenerational transmission of skill. The CTC is accessed frictionlessly by the eligible, consistent with its implementation via the tax system. We represent the spending program implementation method of the CCDF by introducing a stochastic application cost (necessary but not sufficient for receipt), as well as a subsequent rationing stage in which an eligible applicant may be denied access to aid (e.g., never moving off a waiting list). The CCDF application cost, rationing outcome, and child initial skill endowment are realized in the first period of adulthood and are the only source of risk for adults. Consumers subsequently solve a full life cycle problem that includes a retirement phase, during which they collect pensions from Social Security (SS) which are taxed. They cannot save, borrow, or make financial transfers to their children; there is no earnings risk; and the before-tax return per unit of skill in the labor market (the wage rate) is fixed. The government collects revenue via a flat consumption tax and a progressive labor income tax; the average income tax adjusts to balance the government's budget constraint in equilibrium. In the model, the income tax system represents federal taxes excluding transfers from the CTC and the Earned Income Tax Credit (EITC), which are instead modeled as distinct components of disposable income along with the cash benefit component of Temporary Aid for Needy Families (TANF).

Our main experiments compare the expansion of the CTC with that of the CCDF using three measures of welfare changes: first, behind the veil of ignorance at the start of adulthood; second, average changes for new adults conditional on adult skill; and, third, average changes for all adult ages conditional on adult skill. Across steady states, the CCDF expansion generates welfare gains behind the veil of ignorance that are larger than those of the CTC expansion (8.3 versus 2.5 percent of lifetime consumption, respectively). This is due to a larger increase in average adult skill from the CCDF compared to the CTC (6.9 versus less than 0.5 percent, respectively); for both policies, skill inequality also increases because lower percentiles of the skill distribution see relatively smaller gains. Despite the CCDF expansion generating larger increases in average adult skill, the average welfare change for new adults is similar for the CTC and CCDF (1.9 versus 2.1 percent, respectively) in part because the CCDF benefits a smaller share of adults. Average welfare changes for all adults are lower than for new adults, and exhibit small losses after the CTC expansion (-0.2 percent) with net effects of about zero after the CCDF expansion. This occurs because the labor income tax must rise after both expansions to balance the government's budget, which causes welfare losses among older consumers.

Welfare gains from expanding family policies are possible in our model for multiple reasons. For example, children cannot encourage their parents to spend more on skill investment by committing to reimburse these expenses later in life, and adults face life cycle borrowing constraints that prevent them from borrowing against future income to finance consumption when young. There is also uninsurable risk for adults, which stems from stochastic CCDF application costs and child initial skill endowments, as well as CCDF rationing. The government can approximate some of the missing private markets through the tax and transfer system: for example, the CTC increases disposable income for young adults and taxes the income of all adults, shifting resources to earlier ages and addressing life cycle borrowing constraints. Of course, this cash transfer can also be spent on investment in child skill. As for the CCDF, this policy encourages spending on investment in child skill by lowering the expense of a given level of investment for recipients while taxing the income of all adults to recoup the costs; this addresses the inability of children to commit to repay their parents in order to motivate higher skill investment. At the same time, family resources freed up by the childcare subsidy can also be spent on parent consumption or to finance greater leisure when young. Hence, there is conceptual overlap across policies in terms of the missing market they address, which motivates our welfare decomposition exercise.

We offer insight into the market failures each family policy expansion addresses by implementing a decomposition of welfare gains behind the veil of ignorance that builds on the method of Guvenen et al. (2023a). Our decomposition identifies marginal contributions from four components: (i) changes in consumption age profiles; (ii) changes in average child skill outcomes; (iii) changes in average levels of consumption and non-leisure time; (iv) remaining adjustments. Benefits from the CTC expansion stem from redistribution of consumption towards earlier ages (component i), and from changes in the distribution of consumer types and changes in the expected returns to child skill outcomes (component iv).[2] By contrast, benefits from the CCDF expansion are driven by changes in the average levels of consumption and non-leisure time, which are possible because of large increases in average adult skill (component iii). Changes in average child skill outcomes (component ii) contribute positively to welfare gains from both expansions through their effect on the altruism term of parents, but more so for the CCDF. Our results lead us to view the CTC (to a greater extent than the CCDF) as addressing life cycle borrowing constraints by redistributing consumption to earlier ages, while we view the CCDF (to a greater extent than the CTC) as addressing the inability of children and adults to contract with one another in order to promote investment in child skill.

[1] Statistics are in 2016 US dollars and computed as described in Appendix A.1. The cited CCDF spending is for children under 5.

[2] Parental altruism in our model means that parents internalize the expected lifetime utility of their child at the child skill outcome they finance with investment. This object can increase because the child's skill outcome increases (component ii) or because expected lifetime utility goes up at a given level of adult skill (component iv). Although changes in the distribution of quantities across consumers are also present in component (iv) they do not play a large role in this component's contribution towards CTC welfare gains.

Two studies closely related to ours are Guner et al. (2020) and Zhou (2022).[3] Guner et al. (2020) examines a battery of family policies including childcare subsidies and child tax credits. The emphasis is on policy effects on parental labor supply; consistent with this purpose, their analysis accounts for the skill accumulation of parents rather than children. We focus instead on how policies interact with child skill accumulation. We model childcare as an input into skill investment along with parental quality time, representing the time use of children as a primary determinant of their skill accumulation early in life. In this respect we also differ from Zhou (2022), who examines aggregate fertility responses to stylized family policies. Because of our interest in comparing specific policy reforms in the United States, we explicitly model attributes of the CTC that make it a transfer whose level is hump-shaped in household income (Crandall-Hollick, 2018). These attributes are relevant for our policy analysis because changes along these margins are part of the TCJA expansion of the CTC that we simulate in our main policy experiments. Importantly, we also differ from both Guner et al. (2020) and Zhou (2022) in that we incorporate access frictions for subsidized childcare, making explicit an additional distinction between the CTC and CCDF: their implementation via the tax system versus a spending program, respectively.

This paper builds on Moschini (2023) by broadening the set of family policies under consideration to include a child tax credit as found in the US tax code, by partially endogenizing childcare subsidy receipt, by providing a welfare decomposition to aid in identifying and interpreting sources of welfare gains, and (in a supplemental exercise) by examining the transition path dynamics after each policy change. Here, we abstract from heterogeneity in family structure; instead, our focus is on comparing spending-equivalent expansions of two existing policies, rather than comparing expansions of the same policy targeted towards one- versus two-parent families.

For tractability and interpretive clarity our model makes several simplifying assumptions, some of which may impact our results. For example, as noted earlier, in our model framework consumers cannot save, borrow, or make financial transfers to their children, there is no earnings risk, and the pretax return to skill on the labor market is always the same. We discuss the possible impact of relaxing these assumptions in the conclusion.

The paper proceeds as follows. Section 2 overviews the policy environment, Section 3 describes the model, and Section 4 reports model parameterization. Section 5 examines properties of the baseline equilibrium, and Section 6 presents the model experiments and results. Section 7 concludes and suggests directions for future research.

2. The policy environment

This section describes the design, implementation method, and size of the CTC and the CCDF. For each policy, the distributions of spending and recipients for households with children under 5 are presented and discussed. Statistics quoted here are computed as explained in Appendix A.1.

The child tax credit The CTC is a partially-refundable federal tax credit, introduced in 1997 and subsequently expanded in 2001, 2010, 2017, and 2021 (Steuerle, 1990; Crandall-Hollick, 2018, 2021; Goldin and Michelmore, 2022). A tax credit assigns a credit value that depends on taxpayer attributes, such as pretax income or the presence of children; the credit value is then subtracted from the tax liability. If the value of the credit exceeds the tax liability, then a refundable credit rebates at least some of the residual amount to the taxpayer, while a nonrefundable credit does not. Refundable credits may be either fully refundable, so that all of the residual amount is given to the taxpayer, or partially refundable, so that only a portion of the residual is rebated. Altogether, tax credits weakly increase disposable income and are similar to cash transfers. In practice, nonrefundable credits are limited in their ability to deliver transfers to the poor because they are bounded above by tax liabilities (which are low for poor taxpayers under progressive taxation), while fully refundable credits are more effective (Goldin and Michelmore, 2022). Partially refundable credits such as the CTC represent an intermediate case in this regard.[4]

The total CTC per qualifying child in the taxpayer household is the sum of a nonrefundable and a refundable component.[5] To be eligible for any credit, the taxpayer must have at least one dependent child under the age of 17, and taxable income must be within a certain range. As taxable income increases, the total CTC at first increases and then eventually decreases, making it a hump-shaped function of income (Crandall-Hollick, 2018).

The CTC is implemented through the tax system: receipt is possible when an eligible taxpayer correctly completes their income tax form. Such a process identifies eligible households who then face a reduced tax liability and receive the refundable portion of the credit in cash; more than 90 percent of the eligible are estimated to receive CTC transfers (Crandall-Hollick, 2018). Over the 2015-2017 period, an average of $51.5 billion, or 0.27 percent of Gross Domestic Product (GDP), was spent on the nonrefundable and refundable components of the CTC combined, and an average 15 percent of tax filers received transfers from the policy (Internal Revenue Service of the United States, Statistics of Income, 2016, 2017).

[3] Other structural studies of childcare subsidies which allow for endogenous skill accumulation of children include Bastani et al. (2020), Ho and Pavoni (2020), and Daruich (2023). The first two perform a Mirrleesian analysis of optimal childcare subsidies; the former focuses on childcare quality choices, whereas the latter focuses on parental labor supply. The third study focuses on scaling up early childhood interventions such as Head Start rather than analyzing tax credits. Our discussion of missing markets in this paper draws on the analysis of Daruich (2023).

[4] The CTC is similar to the older and better-studied Earned Income Tax Credit (EITC) in that the level of the credit is affected by the presence of children (Marr et al., 2015). The EITC credit level is related to the presence of children in that a much-reduced EITC is available for low-income households without children. Unlike the CTC, the EITC is fully refundable and begins to phase out at lower values of income than the CTC (Crandall-Hollick, 2018; Crandall-Hollick et al., 2021). Because of its attributes, the EITC is received by lower-income families than the CTC.

[5] In US government reports, the nonrefundable component of the tax credit for children is called the "CTC", while the refundable component is referred to as the Additional Child Tax Credit, or "ACTC". In this paper we refer to the sum of these two components as the (total) CTC.

The child care and development fund The CCDF seeks to help working parents with childcare expenses; it began in its current form when the Personal Responsibility and Work Opportunity Reconciliation Act of 1996 consolidated four programs into one (Lynch, 2022).[6] The CCDF gives block grants from the federal government to states; states then use these funds to subsidize childcare expenses for households with children under 13 years of age, although most children receiving the subsidy are under 5 (Chien, 2019a,b, 2020). To qualify to be paid with CCDF funds, childcare providers must meet certain criteria that are intended to keep provider quality comparable to the market average.

Conditional on receipt, the level of the CCDF transfer is the difference between total childcare expenses and the amount contributed by the recipient family (their "copayment", which is increasing in income); thus, the transfer level is decreasing in income and increasing in the quantity of childcare used. Eligibility is set at the state level and must be at least as stringent as federal guidelines, which require that households must contain young children, adults must work or engage in an approved work-related activity, and income must be at or below 85 percent of the state median income for households of a similar size and composition. Most states set a lower income cutoff than the federal guidelines (Chien, 2019a,b, 2020).

The CCDF is implemented via a spending program: to establish eligibility, applicants must contact a social services agency, provide documentation of their work-related activity and income level, and complete an interview with a social worker. Some eligible applicants never receive aid. Therefore, the aggregate amount of CCDF spending and mass of recipients reflects the combined effect of eligibility rules, application rates below 100 percent among the eligible, and rationing conditional on applying (e.g., never moving off of waiting lists). Among the eligible, previous studies have found that those with higher educational attainment (or higher resources) or who face lower application costs are more likely to apply; an estimated 16 percent of eligible applicants do not receive aid, and in practice this rationing outcome is largely uncorrelated with applicant characteristics (Johnson et al., 2011; Guzman, 2019). On an annual basis over the 2015-2017 period, CCDF spending on children under 5 (the majority of CCDF recipients) averaged \$5.5 billion or 0.03 percent of GDP (Chien, 2019a,b, 2020; Administration for Children & Families, 2024a,b,c). Rates of receipt among the eligible are lower than the CTC: estimates indicate that between 10 and 23 percent receive aid.[7] Altogether, about 4 percent of all children under 5 received aid from the CCDF.

Implementation methods and recent policy trends Compared to using the tax system, a spending program introduces more frictions in the process of application and receipt ("access frictions"); the properties of these alternative implementation methods are well-documented, and selecting one or the other is a policy design choice (Steuerle, 1990; Sommartino et al., 2002). In the United States, there has been an increasing trend since the 1980s towards using the tax system to implement social policy objectives, such as transfers to families with children. A case in point of this legislative trend is the 2017 Tax Cuts and Jobs Act, which expanded the CTC.

The distribution of CTC and CCDF spending and recipients Our focus is on how transfers from the CTC and CCDF interact with child skill accumulation during early childhood, when the child is under age 5. For households with children in this age range, we compute the distribution over income quartiles of spending and recipients for the CTC and CCDF. We pool 2015, 2016, and 2017 waves from the Annual Socio-economic Supplement of the Current Population Survey or CPS ASEC (Flood et al., 2024) for data on the CTC, and use CCDF Administrative Data for the 2016 fiscal year for data on the CCDF. Quartile thresholds are estimated using the distribution of income in the CPS ASEC sample of households with children under 5. For further details, see Appendixes A.1.2 and A.1.3.

Fig. 1 displays the distribution of spending (left) and recipients (right) for the CTC and the CCDF (in green and pink, respectively). CTC spending and recipients are concentrated in the second and third income quartiles, while CCDF spending and recipients are concentrated in the first two quartiles. These patterns reflect that the CTC is more generous towards taxpayers with higher tax liabilities and income below an upper bound, while the CCDF offers a subsidy for childcare expenses to applicants with income below an upper bound.[8] Note that the CCDF's frictional implementation method also affects the realized distribution of transfers and recipients for this policy.

In the next section, we specify a model that embeds representations of the CTC and CCDF, as well as the EITC, TANF, and SS pensions.

3. The model environment

The model contains four types of agents: heterogeneous consumers organized into families, the government, a final goods producer, and a childcare provider.

Consumers solve a life cycle problem as illustrated in Fig. 2. Time is discrete with five-year periods and runs forever; consumers live for 75 years and experience 11 periods of adulthood starting at age $j = 1$ when they are 20 years old, prior to which they are a child. Only adults (not children) make decisions. Each consumer becomes the parent of one child at $j = 1$; ages 1 to 4 are the

[6] The rules governing how funding is spent were initially authorized until 2020, after which they have been re-authorized annually by Congress. As for funding, temporary extensions in funding continued from 1996 until 2021, when the American Rescue Plan Act provided permanent annual appropriations.

[7] Together, rationing and application rates of less than 100 percent result in a low rate of aid receipt among the CCDF eligible (the CCDF is not an entitlement, so this pattern is consistent with its mandate). For example, government reports estimate that 15 percent of federally eligible children, or 23 percent of state-eligible children, receive the CCDF (Chien, 2019a,b, 2020). In these reports, eligibility is determined using realized income. An alternative is to define the eligible using potential earnings at some fixed level of labor supply. Using this approach, Guzman (2019) estimates that share of eligible families receiving CCDF aid is around 10 percent.

[8] The small share of CCDF spending that goes to the third quartile is likely due to variations across states in the type of income that counts for eligibility.

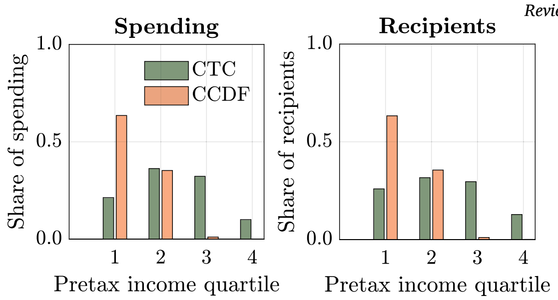

Fig. 1. CTC and CCDF distribution of spending and recipients. **Notes:** Fig. 1 reports the discretized distribution of total spending (left) and recipients (right) by pretax income quartile. For both panels the CTC is in green and the CCDF is in pink. Quartile thresholds are estimated using income of households with children under 5 from the CPS ASEC. Source: CPS ASEC and CCDF administrative data. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

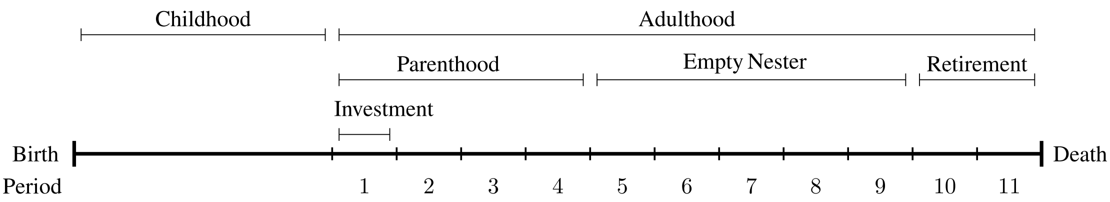

Fig. 2. Phases of the consumer's life cycle.

"parenthood" phase, ages 5 to 9 are the "empty nester" phase, and ages 10 and 11 are the "retirement" phase. The first period of adulthood is a special part of the parenthood phase, the "investment" phase, in which parents invest in their child's skill while their child is under 5 years old.

Adults choose how much to spend on consumption and, before the retirement phase, how many hours of labor to supply. In the first period of adulthood, consumers additionally choose whether to apply for access to a childcare subsidy (the CCDF); application is costly and yields uncertain benefits. They also choose how much to invest in their child's skill using their own time and non-parental childcare time purchased on the market. Skill investment, together with the child's skill endowment, determines the skill level with which the child later begins adulthood. The way investment affects skill is governed by childhood skill accumulation technologies, which parents take as given. Parents are willing to finance costly skill investment expenses because they internalize the effect that skill has on their child's expected lifetime utility in an imperfectly altruistic way. In this model, consumers may not choose to save or borrow, and they cannot make contracts with their children.

At the start of adulthood, consumers are indexed by their adult skill θ_a. The distribution of adult skill at age $j = 1$, $\mu\left(\theta_a\right)$, is endogenous in the sense that it incorporates the skill investment decisions of parents. After the first period of adulthood, skill grows exogenously so that at each age j it is given by the function $\tilde{\theta}\left(j, \theta_a\right)$. At each age of adulthood consumers may differ in their own skill level, but when deciding how much to invest in their child's skill there is additional heterogeneity among adults: adults at age $j = 1$ are also indexed by their child's initial skill endowment, θ_k (mnemonic: k ="kid"), their cost of applying for the CCDF, ξ, their application decision, $\delta \in \{0, 1\}$, and whether or not they are offered the CCDF, $R \in \{0, 1\}$.

During adulthood, consumers may receive transfers from the government which depend on their age (thereby reflecting the presence of children, and whether or not they are retired), whether or not their work, and their level of pretax income (which is entirely made up of earnings). These transfers are allocated by a set of four family policies (the CTC, CCDF, EITC, and TANF) as well as SS pensions; each of the five transfer policies is associated with a vector of policy parameters ($\vec{\pi}_C, \vec{\pi}_N, \vec{\pi}_E, \vec{\pi}_T$, and $\vec{\pi}_P$, respectively). Besides these transfer programs, the government also finances government consumption, set as an exogenous share Θ_G of output. Revenue to finance this spending is raised via a flat consumption tax, τ_c, and a progressive income tax. The income tax collects revenue from pretax income, y, according to $y - \lambda_y y^{1-\tau_y}$, with fixed degree of progressivity given by τ_y that represents the effect of taxes apart from refundable credits. The government adjusts the average income tax, λ_y, to balance its budget. The summary vector $\vec{\pi}_G$ collects all fixed government policy parameters as well as λ_y.

The CCDF is unique among the model's five transfer policies because it is implemented via a spending program, as reflected by two frictions that are encountered in an ordered way during the application process. First, consumers must pay a stochastic cost to apply which does not contribute to government revenue. Second, there is a subsequent rationing stage where applicants may be denied the CCDF with some probability (which is the same for everyone). All other transfer policies are frictionless. The CCDF is also unique among transfer policies in this model in that it is a form of "transfer in kind", in the sense that it is intended to lower the cost of a specific type of expenditure for the consumer, whereas all the other transfers are in cash.

The final goods producer uses a linear production technology that takes efficiency units of labor, H, as its sole input: $Y = H$. The wage rate per efficiency unit of labor supply is therefore equal to 1 and left out of expressions for wages and earnings; in this sense, this model is in partial equilibrium with respect to the wage rate.

The childcare provider offers a perfectly elastic supply of childcare with productivity θ_n at price p. The price of childcare is set as a constant fraction η of the average pre-tax wage of adults aged $j = 1$, so that $p = \eta \int_{\theta_a} \theta_a \mu\left(\theta_a\right) d\theta_a$. This is the only childcare provider in the economy; subsidized childcare and unsubsidized childcare have the same productivity and pre-subsidy price.

The remainder of this section provides more details about the model: Consumer value functions are described in Section 3.1, Section 3.2 describes government transfer policies, and Section 3.3 defines the stationary steady state equilibrium.

3.1. Consumer life cycle problems

3.1.1. The first period of adulthood ($j = 1$)

Each consumer begins adulthood taking their adult skill level θ_a as given. The first period of adulthood is divided into two stages: in stage 1, consumers choose whether or not to apply for the CCDF before they know their rationing outcome or their child's initial skill endowment; in stage 2, the consumer draws their rationing outcome and their child's initial skill, and solves their constrained optimization problem for the investment phase.

Stage 1: Given θ_a, the expected lifetime utility at the start of stage 1 takes an unconditional expectation with respect to ξ, their cost of applying for the CCDF:

$$V_1\left(\theta_a\right) = \mathbb{E}_\xi\left[\tilde{V}\left(\theta_a, \xi\right)\right] \tag{1}$$

Next, the consumer draws ξ from a log-normal distribution that is the same for all consumers: $\log(\xi) \sim \mathbb{N}\left(\mu_N, \sigma_N^2\right)$. Given $\left\{\theta_a, \xi\right\}$, the consumer makes the CCDF application decision $\delta \in \{0,1\}$ by maximizing the expected value in stage 2:

$$\tilde{V}\left(\theta_a, \xi\right) = \max_{\delta \in \{0,1\}} \left\{ \mathbb{E}_{R,\theta_k}\left[V\left(j, \theta_a, \xi, \delta, R, \theta_k\right)\right] \right\} \tag{2}$$

Here, the consumer internalizes that at the start of stage 2 the CCDF rationing outcome R is equal to 1 with probability ϕ_N and zero otherwise, where $R = 1$ indicates that an applicant is granted access to the CCDF. The child's initial skill endowment, θ_k, is drawn from a conditional log-normal distribution where the conditional mean depends on parent skill: $\log\left(\theta_k\right) \sim \mathbb{N}\left(\mu_k + \rho_k\left(\theta_a - \overline{\theta}_a\right), \sigma_k^2\right)$, where $\overline{\theta}_a = \int_{\theta_a} \theta_a \mu\left(\theta_a\right) d\theta_a$.

The outcome of R and θ_k affect the value of applying for the CCDF: if $\delta = 1$ but $R = 0$, then the application cost was incurred with no benefit, whereas if $\delta = 1$ and $R = 1$, then the CCDF will lower the costs of investment and free up resources for consumption at a given level of investment and labor supply (assuming labor supply is chosen to maintain eligibility). At the same time, if $\delta = 0$, then the value of R has no impact on the consumer.

Stage 2: At the start of stage 2, the consumer draws R and θ_k. Given $\left\{\theta_a, \xi, \delta, R, \theta_k\right\}$, consumers solve the investment problem by choosing a level of household consumption, c, labor supply, ℓ, non-parental childcare time, n, and parent quality time with their child, q:

$$\begin{aligned} V\left(j, \theta_a, \xi, \delta, R, \theta_k\right) = \max_{c,\ell,n,q \geq 0}\ & \ln\left(\frac{c}{\Lambda_j}\right) - \psi \frac{(\ell + q)^{1+\frac{1}{\omega}}}{1 + \frac{1}{\omega}} + bV_k\left(\theta_k^a\right) + \beta V\left(j+1, \theta_a\right) \\ & s.t. \\ \left(1 + \tau_c\right) c + pn + \delta\xi \leq\ & y_d\left(y \mid j, \theta_a, \vec{\pi}_G\right) + \delta CCDF\left(y, n \mid j, \xi, R, p, \vec{\pi}_N\right) \\ y =\ & \theta_a \ell \\ \theta_k^a =\ & f\left(n, q \mid \theta_a, \theta_k\right) \end{aligned} \tag{3}$$

where $q + n \leq 1$ and $q + \ell \leq 1$. Here, and in each subsequent period before retirement, labor is restricted to a discrete set of options, $\ell \in \left\{\ell_i\right\}_i$, where each element is between 0 and 1. The age $j = 1$ consumer cares about their own utility, their child's skill outcome, and their own continuation value. For a given level of labor supply, the consumer spends their resources on consumption (which is taxed, and shared with their child according to the equivalence scale Λ_j), childcare expenses, and their CCDF application cost (if $\delta = 1$). Disposable income, $y_d\left(y \mid j, \theta_a, \vec{\pi}_G\right)$, represents the net effect of income taxes as well as the CTC, EITC, and TANF on pretax income; its specific form is provided in equation (9) of Section 3.2. Pretax income y is determined by adult skill and labor supply. In a way that depends on their pretax income (and embeds a working requirement) a CCDF applicant may also be partially reimbursed for childcare expenses from the CCDF if $R = 1$. Finally, the technology $f\left(n, q \mid \theta_a, \theta_k\right)$ governs how investment inputs chosen by the parent combine with the child's initial skill endowment to generate the child's skill level as an adult. We discuss this function and its role in consumer optimization next.

The skill investment decision: costs, benefits, and technology When choosing how much to invest in their child's skill in stage 2 of period 1, parents proceed in two steps. First, they minimize the cost of achieving each feasible skill outcome θ_k^a. Cost minimization weighs the relative prices and productivities of n and q given parent type θ_a and the market's productivity of purchased childcare, θ_n; for a family type, it establishes the cheapest combination of n and q to achieve θ_k^a. Second, parents choose the optimal child skill outcome to balance the marginal costs of raising the child's skill outcome against the marginal benefit. The marginal costs reflect cost-minimization; the marginal benefit arises because parents are altruistic, as per the term $bV_k^a(\cdot)$.

The altruism term is the product of two components: the parent's expectation of their child's lifetime utility at a candidate outcome in stage 1 of $j = 1$, $V_k^a(\cdot)$, and the altruism coefficient, b, which weights the well-being of the child relative to the parent's well-being. Parents are imperfectly altruistic towards their child in the sense that they also care about themselves; for example, setting $b = \beta$ (or equal to 1) does not imply that the objective of the parent and the child coincide perfectly, because children do not care about the utility their parents get from consumption or the disutility they get from non-leisure time. Motivating costly investment via altruism, as we do here (as opposed to specifying a term that gives fixed benefits to a given skill outcome) allows the expected return to the child's skill for the parents to adjust as the environment changes. For example, adult skill may have a different effect on expected lifetime utility if the tax rate rises to finance expanded transfer policies.

Returning to the cost-minimization step of investment optimization, this process takes as a constraint the technology that maps investment inputs to skill outcomes, which in this model consists of two nested CES functions represented by $f\left(n, q \mid \theta_a, \theta_k\right)$. The outermost aggregator is the dynamic function, which combines investment, $I\left(n, q \mid \theta_a\right)$, and the child's initial skill endowment, θ_k, to generate the skill outcome that the child will take into adulthood:

$$f\left(n, q \mid \theta_a, \theta_k\right) = \lambda_\theta \left[(1-\upsilon)\left(\theta_k\right)^{\frac{\chi-1}{\chi}} + \upsilon I\left(n, q \mid \theta_a\right)^{\frac{\chi-1}{\chi}}\right]^{\frac{\chi}{\chi-1}} \tag{4}$$

Here, υ and χ are the CES share parameter for investment and the elasticity of substitution between the child's stock of skill and investment, respectively, while λ_θ scales the output into the space of adult skill. The level of investment is determined by the innermost aggregator, $I\left(n, q \mid \theta_a\right)$, which combines investment inputs n and q within a period:

$$I\left(n, q \mid \theta_a\right) = \left[(1-\gamma)\left(\theta_n n\right)^{\frac{\nu-1}{\nu}} + \gamma\left(\theta_a q\right)^{\frac{\nu-1}{\nu}}\right]^{\frac{\nu}{\nu-1}} \tag{5}$$

In this specification, γ is the CES share parameter on the efficiency units of parent quality time investment input, and ν is the elasticity of substitution between efficiency units of parental time and efficiency units of childcare time. The arguments of the CES aggregator are efficiency units of each input, where θ_a is the productivity of quality time and θ_n is the productivity of non-parental childcare time.

As a thought experiment, one can imagine that given the prices of parental quality time and non-parental childcare time we construct a composite price index in the manner of Moschini (2023). This price index will depend on the prices of q and n and the parameters of the technology that maps investment inputs to the level of investment that they generate. As the price of n falls due to a childcare subsidy, this passes through to the composite price index for investment to an extent mediated by the parameters of (5) and thereby lowers the relative price of investing in child skill compared to other sources of utility. As a result, a substitution effect arises, causing consumers to shift towards the now relatively cheaper source of utility: better child skill outcomes.

As for the dynamic function of equation (4), this technology has the property that higher initial skill endowments have a higher marginal product of investment. To preview an attribute of the model baseline equilibrium, children with higher skill endowments will receive more investment from their parent, all else equal; this pattern is consistent with empirical findings in Datar et al. (2010); Frijters et al. (2013); Grätz and Torche (2016) and Attanasio et al. (2020), among others. When children with high skill endowments are born to parents with low adult skill levels, the level of investment chosen by the parent leaves room for a great deal of improvement in children's skill outcomes—improvement which increased transfers from family policies can instigate.

3.1.2. *Remainder of life cycle ($2 \le j \le 11$)*

After the first period of adulthood, the consumer's level of adult skill is governed by the function $\tilde{\theta}\left(j, \theta_a\right)$, a third-order polynomial given by:

$$\tilde{\theta}\left(j, \theta_a\right) = \left(1 + \beta_1^{age}(j-1) + \beta_2^{age}(j-1)^2 + \beta_3^{age}(j-1)^3\right)\theta_a \tag{6}$$

Note that $\tilde{\theta}\left(1, \theta_a\right) = \theta_a$, consistent with the constraints of the $j = 1$ value function of (3). Although skill grows over the life cycle, because age and θ_a are sufficient to determine what the skill level for any consumer is, the consumer is nevertheless indexed by $\left\{j, \theta_a\right\}$ at ages $2 \le j \le 11$.

Working ages For $2 \le j \le 9$, the consumer solves the problem in (7). They choose consumption and labor supply and continue to split consumption with their child up to and including period $j = 4$, after which the adult lives alone and $\Lambda_j = 1$.

$$\begin{aligned} V\left(j, \theta_a\right) &= \max_{c, \ell \ge 0} \ \ln\left(\frac{c}{\Lambda_j}\right) - \psi \frac{\ell^{1+\frac{1}{\omega}}}{1+\frac{1}{\omega}} + \beta V\left(j+1, \theta_a\right) \\ &\quad s.t. \\ \left(1+\tau_c\right) c &\le y_d\left(y \mid j, \theta_a, \vec{\pi}_G\right) \\ y &= \tilde{\theta}\left(j, \theta_a\right)\ell \end{aligned} \tag{7}$$

Retirement phase When $10 \le j \le 11$, the consumer solves the problem in (8), in which there is no disutility of labor and they earn no pretax income because they do not work, and the continuation value drops out of the objective function in the last period of life.

$$V\left(j,\theta_a\right)= \ln\left(\frac{c}{\Lambda_j}\right)+\mathbb{I}_{j<11}\beta V\left(j+1,\theta_a\right) \tag{8}$$

$$s.t.$$

$$\left(1+\tau_c\right)c \le y_d\left(0\mid j,\theta_a,\vec{\pi}_G\right)$$

Because there is no borrowing or saving in this model, the consumer's disposable income during retirement is equal to their after-tax SS pension, which is embedded in y_d.

3.2. Government policies

Section 3.2.1 defines disposable income; in addition, the functional forms and policy parameter vectors of the CTC and CCDF are defined and qualitative illustrations of each transfer as a function of pretax income are provided for intuition. Because the transfer value at a given level of income depends on parameterization, these figures are stylized and do not include dollar values. We then outline the structure of the three remaining transfer policies, the EITC, TANF, and SS, whose specific functional forms and policy parameter vectors are provided in Appendixes A.2.4, A.2.5, and A.2.6, respectively. Section 3.2.2 summarizes the work- and age-eligibility requirements of the five transfer programs in our model. The extent of means-testing (i.e., each transfer policy's income eligibility rules) is a consequence of specific parameterizations and its discussion is deferred until the model parameterization of Section 4. The government's budget constraint is defined in Appendix B.2.

3.2.1. Government tax and transfer policies

Disposable income is a function that conditions on whether or not the consumer works, which in our model is equivalent to whether or not y is positive.

$$y_d\left(y\mid j,\theta_a,\vec{\pi}_G\right)=\begin{cases}\lambda_y y^{1-\tau_y}+CTC\left(y\mid j,\lambda_y,\tau_y,\vec{\pi}_C\right)+EITC\left(y\mid j,\vec{\pi}_E\right) & \text{if } y>0\\ TANF\left(j,\vec{\pi}_T\right)+SS\left(j,\theta_a,\lambda_y,\tau_y,\vec{\pi}_P\right) & \text{if } y=0\end{cases} \tag{9}$$

When $y>0$, the first term on the right-hand side of equation (9) represents after-tax income and takes its functional form from Heathcote et al. (2017). The second and third terms are the CTC and the EITC, which depend on the earned pretax income, the age of the consumer, and the relevant vector of policy parameters. Note that the CTC is also affected by the income tax liability governed by λ_y and τ_y. In the second line, when $y=0$, the consumer receives either TANF or an after-tax SS pension: TANF is available to consumers with young children, and SS pensions are available to retirees. If the consumer is not eligible for a transfer as represented in equation (9) due to age or income level, then the value of that transfer is set to zero.

It is important to note that the tax parameters τ_y and λ_y are independent of the age of the consumer and represent the effect of the income tax system before incorporating the CTC and EITC, which are refundable tax credits whose receipt or intensity is tied to the presence of children.[9] In the model parameterization of Section 4, we explain how we assign a value to τ_y by estimating the progressivity of the tax system before refundable credits are applied.

Next, we explain the five transfer policies that the government operates. We begin with the CCDF, which is not included in the definition of disposable income and instead enters the budget constraint of the consumer in (3) as a separate term.

The child care and development fund We represent the CCDF as a proportional subsidy for market-based childcare for families with young children; our specification captures the fact that only lower-income families are eligible and that poorer recipients get more aid at a fixed level of childcare use. It also captures the fact that the level of the CCDF transfer received is increasing in the quantity of childcare used. The total value of the CCDF transfer is:

$$CCDF\left(y,n\mid j,\xi,R,p,\vec{\pi}_N\right)=\begin{cases}pnR\mathbb{I}_{y\in(0,\overline{y}_N)}\tau_N\left(y\mid\vec{\pi}_N\right)+\min\left\{\kappa_N,\xi\right\} & \text{if } j=1\\ 0 & \text{otherwise}\end{cases} \tag{10}$$

This transfer only takes a positive value for families with children age 5, so that $j=1$; if this condition is met, then the transfer is the sum of a proportion of childcare expenses (first term) and a reimbursement for the application cost (second term).

Childcare expenses are determined by the price of childcare, p, and the quantity of childcare purchased, n; subsidized care providers are not different from unsubsidized ones in our model environment. If $R=1$ and the income of the family is above 0 (i.e., the parent works) and below an eligibility threshold $\overline{y}_N$, then the government pays a portion of childcare expenses, determined as described in the next paragraph. The income eligibility threshold is set to a percentile Θ_N of the income distribution for parents raising children under age 10 and adjusts in equilibrium. This reflects that the income eligibility requirement of the CCDF are related to median income for similar families in practice, where "similar" families are parents with children aged under the statutory eligibility cutoff of 13 (here, we approximate with the closest age threshold of 10 given the model's 5-year periods). Note, however, that in this model childcare is only used before the child is 5 years old, and therefore this is the age range where the CCDF can be received in our model.

[9] Any completely nonrefundable tax credits, for example the Child Care and Dependent Tax Credit, are left embedded in the estimation of τ_y. Note that the size of this credit is small.

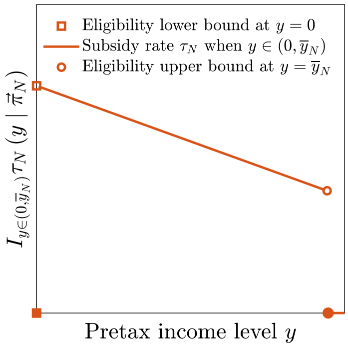

Fig. 3. Stylized illustration of the CCDF subsidy rate including income eligibility requirements. **Notes:** Fig. 3 illustrates the subsidy rate of the CCDF as a function of income for an arbitrary set of parameters, holding fixed median income $\overline{y}_{pop}$. Because this illustration is pre-parameterization and intended to be qualitative, the x- and y-axis values are not labeled except for the point where both values are zero.

The CCDF subsidy rate, $\tau_n\left(y \mid \vec{\pi}_N\right)$, is a linear function of pretax income normalized by median income $\overline{y}_{pop}$:

$$\tau_N\left(y \mid \vec{\pi}_N\right) = \max\left\{\beta_{N,0} + \beta_{N,1}\left(\frac{y}{\overline{y}_{pop}}\right), 0\right\} \tag{11}$$

where the subsidy rate is bounded below by zero. The linear relationship with income represented in equation (11) is how our model allows for the subsidy rate to be decreasing in income and increasing in the number of hours of childcare used, rather than literally modeling the family copayment. We normalize with respect to the median income of the population so that, as the economy grows, what matters for the CCDF subsidy rate is how poor one is relative to the median rather than how poor one is in absolute terms. Across equilibria, one needs the same change in income as a percent of median income in order to get the same change in the subsidy rate; without this attribute, if the economy grew by enough, only those with income in the lowest part of the income distribution would get a positive subsidy rate, while the higher-income eligible might instead be assigned a subsidy rate of zero even if their income is below $\overline{y}_N$.

Turning to the second term in equation (10) when $j = 1$, an CCDF applicant receives a reimbursement (sometimes referred to as an application cost subsidy) of up to κ_N of the application cost, regardless of the CCDF rationing outcome. This reimbursement is not affected by income, and is normalized to zero in the baseline economy. We vary this parameter in our main counterfactual experiment when we expand the CCDF; changing the application cost reimbursement represents a reduction in the frictions faced when seeking to access CCDF aid. Note that, because the amount of this reimbursement is never more than the cost of applying, this component of the CCDF is never a net-positive cash transfer to the applicant family; we therefore view a nonzero value for κ_N as not changing the transfer-in-kind attribute of the CCDF.

Equations (10) and (11), together with the rationing rate and the distribution of application costs, imply that the vector of CCDF parameters is: $\vec{\pi}_N = \left\{\Theta_N, \beta_{N,0}, \beta_{N,1}, \kappa_N, \phi_N, \mu_N, \sigma_N^2\right\}$. Fig. 3 illustrates the linear subsidy rate of the CCDF as a function of income for an arbitrary set of parameters, holding fixed median income $\overline{y}_{pop}$. Specifically, the term being graphed is $\mathbb{I}_{y\in\left(0,\overline{y}_N\right)}\tau_N\left(y \mid \vec{\pi}_N\right)$ from equation (10), using the functional form of equation (11). This term incorporates the upper and lower bounds on pretax income for eligibility, shown in the subscript of the term $\mathbb{I}_{y\in\left(0,\overline{y}_N\right)}$, where the lower bound is equivalent to a work requirement because it imposes that earnings be greater than zero. In the figure, the subsidy rate is set to zero at either end of the range of eligible incomes, as indicated by with the hollow square and circle at either end of the sloped line. For eligible incomes $y \in \left(0, \overline{y}_N\right)$, the subsidy rate $\tau_N\left(y \mid \vec{\pi}_N\right)$ is shown to be decreasing in income ($\beta_{N,1} < 0$) and to have a nonzero y-intercept ($\beta_{N,0} > 0$). Scaling Fig. 3's subsidy rate by the family's childcare expenses would yield the transfer level that the family receives from the CCDF. The term $\min\left\{\kappa_N, \xi\right\}$ from equation (10) is not graphically illustrated because it is normalized to zero in the baseline economy

The child tax credit Our model representation of the CTC follows the statutory structure of this tax credit. The CTC only takes positive values when j is between 1 and 3, so that the child under age 15 (the closest age threshold to 17, the statutory upper bound on child age, given 5-year periods). The credit level per child is calculated as the sum of a nonrefundable component and a refundable component:

$$CTC\left(y \mid j, \lambda_y, \tau_y, \vec{\pi}_C\right) = \begin{cases} CTC_n\left(y \mid \lambda_y, \tau_y, \vec{\pi}_C\right) + CTC_r\left(y \mid \lambda_y, \tau_y, \vec{\pi}_C\right) & \text{if } j = 1, 2, \text{ or } 3 \\ 0 & \text{otherwise} \end{cases} \tag{12}$$

Both components depend on pre-tax income; because this is a partially refundable credit, the income tax liability will affect its level, and therefore λ_y and τ_y appear as arguments in equation (12). The non-refundable component is equal to the lower bound of two constraints and is required to be nonnegative:

$$CTC_n\left(y \mid \lambda_y, \tau_y, \vec{\pi}_C\right) = \max\left\{\min\left(\overline{\kappa}_{C,n} - \beta_{C,n}\left(y - \overline{y}_{C,n}\right)\mathbb{I}_{y>\overline{y}_{C,n}}, y - \lambda_y y^{1-\tau_y}\right), 0\right\} \tag{13}$$

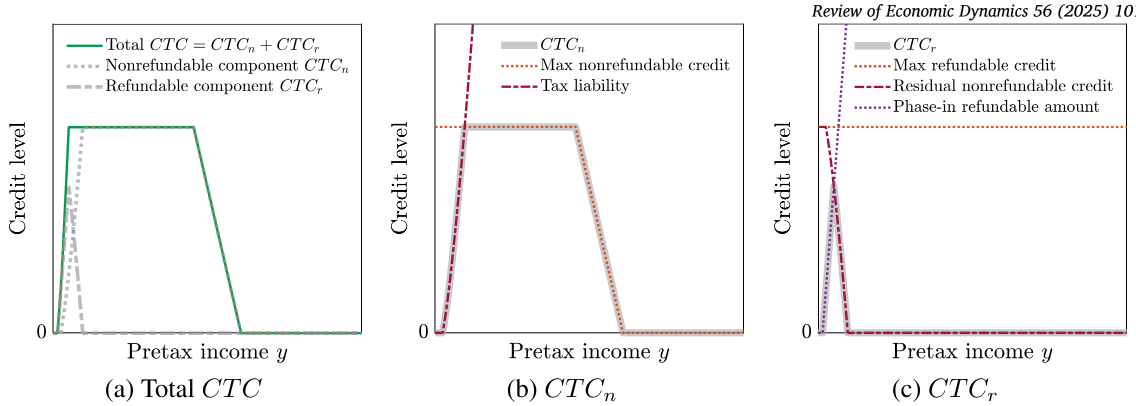

Fig. 4. Stylized illustration of the total CTC and its two components. **Notes:** Fig. 4 illustrates the total CTC of equation (12) in Fig. 4a (solid line, with its two components as dashed lines). The nonrefundable component of equation (13) is shown in Fig. 4b, and the refundable component of equation (14) in Fig. 4c. Each component (thick grey line) is the minimum of a set of constraints (dashed lines), which are listed in the legend in the same order as in the component's equation in the main text. Additionally, each component is bounded below by zero. Because this illustration is pre-parameterization and intended to be qualitative, the x- and y-axis values are not labeled except for the point where both values are zero.

Moving from left to right on the right-hand side of equation (13), the first constraint is the maximum nonrefundable amount. This amount begins at a value $\overline{\kappa}_{C,n}$ and phases out at a rate $\beta_{C,n}$ for every additional dollar of earnings once income exceeds the phaseout threshold, $\overline{y}_{C,n}$. The second constraint is the consumer's income tax liability before tax credits. If the minimum of these two constraints is negative, then the nonrefundable component of the CTC is set to zero.

Here, the maximum credit level and phaseout threshold are set as constant fractions of median income, $\overline{y}_{pop}$, so that their levels in units of consumption adjust in equilibrium. Specifically, $\overline{\kappa}_{C,n} \equiv \Theta_{C,\kappa,n}\overline{y}_{pop}$ and $\overline{y}_{C,n} \equiv \Theta_{C,y,n}\overline{y}_{pop}$.

The refundable component of the tax credit is equal to the lower bound of three constraints and is also required to be nonnegative:

$$CTC_r\left(y \mid \lambda_y, \tau_y, \vec{\pi}_C\right) = \max\left\{\min\left(\overline{\kappa}_{C,r}, \overline{\kappa}_{C,n} - CTC_n\left(y \mid \lambda_y, \tau_y, \vec{\pi}_C\right), \beta_{C,r}\left(y - \overline{y}_{C,r}\right)\right), 0\right\} \tag{14}$$

Moving from left to right on the right-hand side of equation (14), the first constraint is the level of the maximum refundable credit, $\overline{\kappa}_{C,r}$; the second constraint is the difference between the maximum nonrefundable credit and the realized value of the nonrefundable component; and the third constraint is the refundable amount, set to the refundability rate, $\beta_{C,r}$, multiplied by the difference between pretax income and the refundability threshold, $\overline{y}_{C,r}$. If the minimum of these three constraints is negative, then the refundable component of the CTC is set to zero. Similar to the non-refundable component, the maximum refundable credit and refundability threshold are parameterized as constant fractions of median income, so that $\overline{\kappa}_{C,r} \equiv \Theta_{C,\kappa,r}\overline{y}_{pop}$ and $\overline{y}_{C,r} \equiv \Theta_{C,y,r}\overline{y}_{pop}$.

The vector of CTC parameters is given by: $\vec{\pi}_C = \left\{\beta_{C,n}, \beta_{C,r}, \Theta_{C,\kappa,n}, \Theta_{C,\kappa,r}, \Theta_{C,y,n}, \Theta_{C,y,r}\right\}$. This vector combines parameters that appear in equations (12), (13), and (14). To illustrate the form of the CTC, Fig. 4 displays the total CTC transfer and its two components, each as a function of income. Specifically, Fig. 4a shows the total tax credit of equation (12) (solid line), which is the sum of its two components (dashed lines); Fig. 4b illustrates the nonrefundable component of equation (13); and, Fig. 4c illustrates the refundable component of equation (14). Each component (thick grey line) is the minimum of a set of constraints (dashed lines), which are listed in the Figure's legend in the same order as in the component's equation in the text. Additionally, each component is bounded below by zero.

Remaining transfer policies: EITC, TANF, and SS The EITC is similar to the CTC in that it phases in to a maximum credit value and then phases out as earned income increases; however, it differs structurally in three main ways. First, the EITC is fully refundable and starts to phase in with the first dollar of earnings, so that consumers receive the full value of the credit regardless of their tax liability as soon as they start working; second, it is available when $1 \le j \le 4$, which includes the last period of the parenting phase; and, third, it is available in a reduced form for working adults whose children have left home (when $4 < j \le 9$). As with the CTC, the maximum credit level and phaseout income threshold of the EITC are parameterized as fractions of median income in equilibrium. For further details, see Appendix A.2.4.

TANF is modeled as a cash benefits available to non-working parents age $j = 1$ (that is, with children under 5). The level of the transfer for TANF is set as a constant fraction of median income in equilibrium. For further details, see Appendix A.2.5.

Lastly, SS pensions are indexed by the adult skill of the consumer; their after-tax value is determined by a pretax benefit that replaces the average earnings of the skill type in equilibrium at a fixed rate, ζ_P, the share of the pretax benefit that is taxed, and the income tax charged by the government on the taxed portion of the benefit which is governed by λ_y and τ_y. Pensions are available during the retirement phase when $10 \le j \le 11$. For further details, see Appendix A.2.6.

3.2.2. *Eligibility for transfers in the model by work status and age*

Before parameterization, we can make statements summarizing the work and age eligibility rules of the model's five transfer policies: the CTC, CCDF, EITC, TANF, and SS pensions. TANF and SS are only received by non-working adults (that is, with zero pretax income); by contrast, the CTC, EITC, and CCDF require that consumers work (that is, have positive pretax income) to receive a

$j=1$	$j=2$	$j=3$	$j=4$	$j \in [5,9]$	$j \in [10,11]$
CCDF					
TANF					
CTC	CTC	CTC			
EITC	EITC	EITC	EITC	EITC*	SS

Fig. 5. Transfer policy age eligibility. **Notes:** Fig. 5 lists, for each age j, the transfers from which a consumer might receive aid, assuming work eligibility. EITC* denotes the reduced EITC available to adults without children at home.

transfer. As for age-eligibility, Fig. 5 lists the set of transfers that a consumer could receive at each age j, assuming that their working status makes them eligible. Conditional on being work- and age-eligible, the level of the transfer as a function of recipient attributes depends on the consumer's realized income and the specific parameters of each policy. Although not reflected in this figure, y could be such that the transfer takes a value of zero even if the consumer is work- and age-eligible.

3.3. Stationary steady state equilibrium

Given a conditional distribution of initial skill endowments for each adult skill, a goods production technology, a childcare pricing rule, and a set of government policy vectors and tax parameters, a stationary steady state equilibrium is agent choices and objective functions, consumer expectations about lifetime utility at each skill level, $V_k^a(\cdot)$, a skill distribution among new adults, $\mu\left(\theta_a\right)$, and an average income tax rate, λ_y, such that agents optimize, expectations are rational, the adult skill distribution is stationary, and the government balances its budget constraint. An extended equilibrium definition is provided in Appendix B.3.

4. Model parameterization

The model parameters can be separated into those estimated outside of the model, reported in Table 1, and those estimated inside of the model, reported in Table 2.

Panel A of Table 1 reports parameters related to government taxation and consumption as well as consumer demographics and preferences. Unless otherwise noted, we parameterize to the years 2015-2017. The consumption tax rate, τ_c, is estimated following the method of Mendoza et al. (1994) using Organization of Economic Cooperation and Development (OECD) data as described in Appendix A.2.1. The income tax progressivity parameter, τ_y, is estimated using the CPS ASEC as described in Appendix A.2.2. Crucially, we estimate the progressivity of the model's tax system to reflect the effect of the federal tax system excluding the CTC and EITC; these refundable credits, along with other family policy transfers, are modeled as separate objects in the consumer's budget constraint. Government consumption as a share of output, Θ_G, is estimated using data from the US Bureau of Economic Analysis (BEA) reported in BEA Table 1.1.10 (2024). The elasticity of non-leisure time, ω, is set to 1.5, the midpoint of range of Frisch elasticity values used in the macroeconomic literature reported in Keane and Rogerson (2015). The vector of age-specific consumption-equivalence scales, $\left\{\Lambda_j\right\}_j$, accounts for the presence of a child following the 1994 OECD modified scales based on the findings of Hagenaars et al. (1994). The unconditional average of the log child skill endowments, μ_k, is normalized to 0. The discount rate, β, is set to be consistent with an annual value of 0.96 and the model's period length of five years. Lastly, the discrete set of labor supply options, $\left\{\ell_i\right\}_i$, is assumed to take values at 8-hour intervals from 0 to 48 percent of the time endowment (which corresponds to 100 hours per week).

Panel B of Table 1 contains values for the parameters governing the CTC for the 2016 Fiscal Year, set using values contained in the Congressional Research Service (CRS) report by Crandall-Hollick (2018). Dollar amounts are for one child, converted into per-parent terms and then averaged across married and single households using shares from Table 1 of Moschini (2023). Next, dollar values are expressed as ratios of median household income. This normalization divides by a value of $57,657 in 2016 US dollars, which is the average of annual median values over the 2015-2016 period from American Community Survey briefs published by the United States Census Bureau (Guzman, 2017, 2018), converted to 2016 dollars using the Consumer Price Index or CPI (Bureau of Labor Statistics, US Department of Labor, 2024). Rates are averaged across single and married filing jointly values, using the same population shares as used for dollar amounts. Overall, the parameterized CTC targets working adults with children under 15 who are closer to the middle of the income distribution.

Panel C contains the CCDF policy parameters. The eligibility threshold percentile is set to 27 to match the across-state average computed using Government Accountability Office (GAO) reports Chien (2019a,b, 2020), and the parameters governing the childcare subsidy rate are estimated using CCDF administrative data. Both estimations are described in more detail in Appendix A.2.3. The government reimbursement for the CCDF application cost is normalized to zero in the baseline. The share of applicants offered aid (as opposed to being rationed out) is set to the national estimate using data from the National Survey for America's Families (NSAF) found in Guzman (2019), while we estimate the variance of CCDF application costs using across-state variation in application rates found in the same publication. The remaining CCDF policy parameter (the average application cost), is not shown in Panel C because it is calibrated internally. Overall, note that the parameterized transfer policies in our model are such that the CCDF targets working adults with young children who are relatively poor.

Panel D contains the three remaining policies: EITC, TANF, and SS. The policy parameters for the EITC are set using information from the CRS report Crandall-Hollick et al. (2021); parameters for TANF are set using a report from the Administration for Children

Table 1
Externally estimated parameters.

Symbol	Parameter description	Data source	Parameter value
Panel A: Government tax and consumption policy, demographics, preferences			
τ_c	Consumption tax rate	OECD	0.043
τ_y	Income tax progressivity	CPS ASEC	0.061
Θ_G	Government consumption (share output)	BEA	0.175
ω	Elasticity of non-leisure time	Assumption	1.50
$\{\Lambda_j\}_j$	Consumption-equivalence scales	OECD	$1+\mathbb{1}_{j\leq 4}0.3$
μ_k	Unconditional average log (θ_k)	Normalization	0
β	Patience	Assumption	0.815
$\{\ell_i\}_i$	Labor supply choice set	Assumption	$\{0.08\,(i-1)\}_{i=1}^{7}$
Panel B: CTC policy $\vec{\pi}_C$			
$\Theta_{C,\kappa,n}$	Maximum non-refundable credit ratio	CRS and US Census, normalized*	$1,000
$\Theta_{C,\kappa,r}$	Maximum refundable credit ratio		$1,000
$\Theta_{C,y,r}$	Refundability threshold ratio		$1,815
$\Theta_{C,y,n}$	Phase-out threshold ratio		$59,200
$\beta_{C,n}, \beta_{C,r}$	Phase-out and phase-in rates		(0.150, 0.050)
Panel C: CCDF policy $\vec{\pi}_N$			
Θ_N	Eligibility threshold percentile	GAO	27
$\beta_{N,0}, \beta_{N,1}$	Childcare subsidy rate	CCDF administrative data	(0.957, −0.348)
κ_N	Application cost reimbursement	Normalization	0
ϕ_N	Rationing: share recipients | applied	NSAF	0.840
σ_N^2	Variance of application cost distribution	NSAF	0.114
Panel D: EITC, TANF, and SS policies			
$\vec{\pi}_E$	EITC policy vector	CRS	See Appendix A.2.4
$\vec{\pi}_T$	TANF policy vector	ACF	See Appendix A.2.5
$\vec{\pi}_P$	SS policy vector	BEA and SSA	See Appendix A.2.6
Panel E: Skill accumulation			
Early childhood			
χ	Elasticity of substitution I and θ_k	ECLS-B	0.380
υ	Share on I		0.490
ν	Elasticity of substitution q and n		0.559
γ	Share on q		0.324
Adulthood			
$\beta_1^{age}, \beta_2^{age}, \beta_3^{age}$	Life cycle wage profile	US Census	(0.509, −0.090, 0.005)

Notes: Table 1 reports externally estimated parameter symbols, descriptions, data sources, and values. * Panel B dollar values are normalized by US median household income, computed as described in the text.

and Families (ACF), a subsidiary of the US Department of Health and Human Services (Office of Family Assistance, 2015). Finally, parameters for the SS pension program are set using information found in Social Security Administration (SSA) reports (Social Security Administration of the United States, 2016b,a) and tabulations from US Bureau of Economic Analysis (2023b,a). The specific parameters of the EITC, TANF, and SS are reported in their respective appendix; qualitatively, the EITC targets working adults who are relatively poor (in contrast to the CTC), while TANF targets non-working adults when they are young, and SS pensions are restricted to non-working adults during the retirement phase, so that both TANF and SS pensions are targeted towards the poor.

Finally, in Panel E, we report parameters related to skill accumulation during early childhood and adulthood. Beginning with the technologies governing skill accumulation during early childhood, we set these parameters using estimates based on data from the Early Childhood Longitudinal Survey - Birth Cohort (ECLS-B) from Moschini (2023), as described in Appendix A.2.7. We parameterize the third-order polynomial that governs the age profile of skill (which is the same as wages in our framework), $\tilde{\theta}\left(j,\theta_a\right)$, using a calibration that targets wage growth moments estimated using US Census data in Lagakos et al. (2018) as explained in Appendix A.2.8.

Qualitatively, the technology parameter values for early childhood in Panel E indicate that investment technology exhibits CES complementarity between investment inputs (governed by ν), as well as between investment and the current stock of child skill (governed by χ). Given the functional form that we assume, the marginal product of investment is increasing in the level of initial skill for any value of χ; however, the elasticity of the marginal product of investment does depend on χ and is increasing in the initial skill of the child when $\chi < 1$, as is the case in our parameterization. In our model, the higher the child's initial skill endowment, the larger the marginal product of investment during early childhood; the extent to which high- and low-skill parents face high returns to investment depends on the correlation of the initial child skill endowment with the parent's skill. The CES share on investment, υ, is 0.49, indicating that both investment and the initial skill endowment of the child play a large role in determining the child's skill outcome. The CES share on parental quality time, γ, indicates that both parental and non-parental childcare time play a role in generating investment. This means that subsidies to childcare's price have the potential to reduce investment's composite price index significantly, thereby motivating a substitution towards skill investment by the household and a consequent improvement in skill outcomes for the child.

Table 2
Internally estimated parameters.

Category	Sym.	Parameter description	Value	Moment description	Data	Model
Panel A: Jointly calibrated	ψ	Marginal disutility non-leisure time	6.71	Age 25-55 average hours labor ℓ	0.31	0.31
	λ_θ	Skill scaling factor	4.83	Age 20 average hours childcare n	0.31	0.31
	b	Altruism coefficient (share β)	0.33	Age 20 Corr$\left(\theta_k^{a*}, y^*\right)$	0.32	0.32
	μ_N	Ave. app. cost (share $\overline{y}_{pop}$)	0.14	Age 20 share receiving CCDF	0.04	0.04
	σ_k^2	Variance of log $\left(\theta_k\right)$	1.03	Age 25-34 log (y) p50-p10	1.39	1.39
	ρ_k	Intergenerational persistence of skill	0.09	Age 20 Corr$\left(\theta_k, \theta_a\right)$	0.06	0.06
	ζ_P	SS replacement rate	0.20	SS spending (share output)	0.05	0.05
Panel B: Set proportionally	θ_n	Childcare productivity	0.70	Age 20 workers ave. log prod. ratio	0.00	0.00
	η	Childcare pre-policy price ratio	0.15	Age 20 workers price ratio	0.18	0.18

Notes: Table 2 presents internally estimated parameters. Panel A contains parameters calibrated jointly inside the model; Panel B contains parameters set proportional to endogenous objects in the baseline equilibrium.

Moving to parameters whose values are estimated inside the model, Panel A of Table 2 presents seven parameters that are calibrated jointly to bring seven model moments as close as possible to their empirical counterparts: the marginal disutility of non-leisure time, ψ, the skill scaling factor, λ_θ, the degree of altruism, b, the mean childcare subsidy application cost, μ_N, the variance of the initial child skill distribution, σ_k^2, the intergenerational persistence of mean skill, ρ_k, and the SS pension replace rate, ζ_P. Although all of these parameters may affect any of the target moments to some extent, in Panel A parameters are presented in the same row as the moment used to update them during the calibration procedure.

The first moment in Panel A is computed for adults aged between 25 and 55 using the CPS ASEC. The second moment averages values from Moschini (2023) across family structures using population weights, while the third moment is directly reported in the same study. The fourth is computed using information on the number of households with age-eligible children and the number of households receiving CCDF funds for 2015, 2016, and 2017, taken from government reports (Chien, 2019a,b, 2020). The fifth target moment is the difference between the 50th and 10th percentiles of the logged pretax income distribution for the United States, averaged over its annual values for 2015, 2016, and 2017, using values from the Global Repository of Income Dynamics database Guvenen et al. (2023b).[10] Next is the correlation between maternal cognitive skill and child cognitive skill when the child is 1 year old, set to the value reported in the online appendix of Cunha et al. (2010). Finally, we discipline the replacement rate for SS pensions using tabulations from the National Income and Product Accounts (NIPA), specifically BEA Table 1.1.5 (2022a) and BEA Table 2.1 (2022c). Overall, the calibrated model aligns well with the data.

We choose the target moments in Panel A of Table 2 because we want the model baseline equilibrium to reflect the following broad attributes of reality (with the relevant target moment provided in parentheses). First, the pass-through of skill to output for prime working-age adults (average labor supply, ages 25-55). Second, the intensity of childcare usage (average hours of childcare time for children under 5). Third, the degree of intergenerational mobility, reflecting the payoff-relevance of being born to one kind of parent versus another in terms of the effect it has on expected well-being during adulthood (the correlation of child skill outcomes at age 5 and family income).[11] Fourth, the extent of access to CCDF aid (the share of families with young children receiving the CCDF). Fifth, the extent of heterogeneity among poor young adults who are potentially eligible for family transfers (the p50-p10 log income difference, ages 25-34). Sixth, the degree to which children with high potential skill outcomes are present among the poor (the correlation of initial child skill and parent adult skill). Finally, the amount of resources spent on SS pensions (share of output spent on SS). Appendix C.1 provides additional details on the internal calibration; in Appendixes C.5.1 and C.5.2 we perform robustness exercises where we change the set of target moments and examine the impact on our main results.

Panel B of Table 2 reports parameters whose values are set proportional to other equilibrium quantities in the baseline and held fixed in the counterfactual experiments. The childcare productivity, θ_n, is calibrated so that the expected logged ratio of parental and nonparental childcare productivities is equal to zero for working parents, to be consistent with the identification of skill technology parameters. The pre-subsidy price ratio between non-parental and parental time, η, is chosen so that in the baseline equilibrium the average post-subsidy and after-tax price ratio of investment inputs chosen by working parents with children under 5 matches what is observed in the data as reported in Table 2 of Moschini (2023).

5. Properties of the baseline equilibrium

5.1. Selection into CCDF application and receipt

Fig. 6 displays the within-cell percentage who are eligible for, apply for, and receive CCDF aid, broken down by quartiles of application cost and adult skill in Fig. 6a and by quartiles of child initial skill endowment and adult skill in Fig. 6b. Comparing patterns of application and receipt in our baseline equilibrium with the findings of previous studies is a useful way to examine whether

[10] This method of disciplining properties of the initial skill distribution using moments on earnings inequality is based on Huggett et al. (2006).

[11] We report a robustness exercise for our main results in Appendix C.5.2 in which we target an alternative statistic to measure intergenerational mobility, the rank-rank correlation of family income for the core sample of Table 1 from Chetty et al. (2014). Our main results are unchanged.

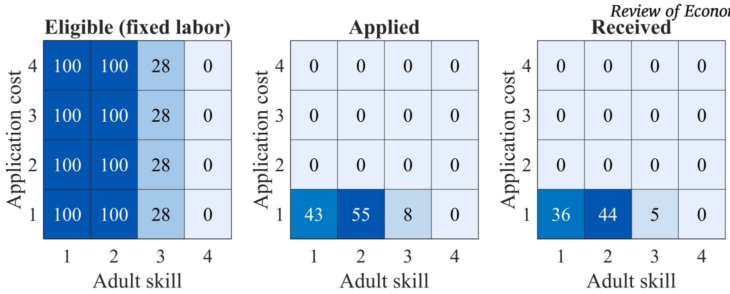

(a) By application cost and adult skill

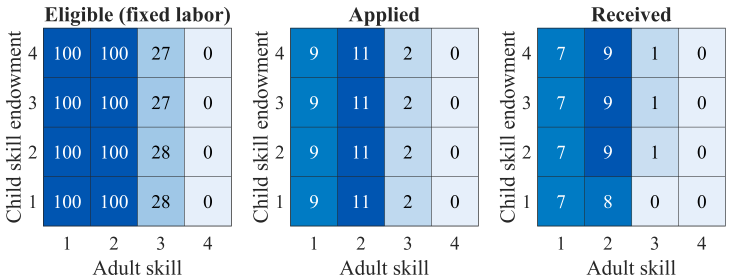

(b) By child initial skill endowment and adult skill

Fig. 6. Selection into CCDF application and receipt. **Notes:** Fig. 6 displays the within-cell percentage who are eligible for, applying for, and receiving CCDF aid for two discretized joint distributions. Fig. 6a displays rates by CCDF application cost and adult skill; Fig. 6b displays rates by child initial skill endowment and adult skill. Quartiles are assigned using the unconditional distributions in the baseline equilibrium.

consumers in our model exhibit reasonable patterns of sorting into CCDF receipt. Here, following the method of Guzman (2019), we define the eligible by comparing potential income at a fixed level of labor supply (16 hours) to the income eligibility threshold of the CCDF in the baseline equilibrium; subsidy receipt requires application, being offered aid, and that the consumer choose a level of labor supply that makes them eligible for a positive transfer.[12] Quartiles are assigned using unconditional distributions for each variable in the baseline equilibrium.

Among the eligible, Fig. 6a indicates that CCDF applicants and recipients tend to have low application costs; given application costs, they also tend to have relatively higher adult skill among the eligible. Fig. 6b indicates that application and receipt are not strongly related to the child's initial skill endowment; this weak relationship is a consequence of CCDF application decisions occurring before the draw of the child's initial skill endowment. This timing matters: parents cannot precisely predict their child's initial skill when deciding to apply, because its correlation with their own adult skill is low. However, a low degree of selection into CCDF receipt by θ_k does arise after the application decision. This is due to adults who, after their draws of R and θ_k are realized, choose a level of labor supply too high to qualify for CCDF aid. This situation arises among applicants with a low θ_k draw, for whom it is not worthwhile to distort labor supply downward in order to be eligible for CCDF aid; it is not worthwhile for such applicants because the optimal level of investment is increasing in θ_k, as demonstrated in Figure C.2a of Appendix C.2.

The patterns of selection into CCDF application and receipt shown in Fig. 6 are qualitatively consistent with the results of Johnson et al. (2011), which indicate that those receiving the CCDF have relatively more resources compared to the pool of eligible non-recipients (higher θ_a), and the results of Guzman (2019), which indicate that higher-educated parents (higher θ_a) and those facing lower application costs (lower ξ) are more likely to apply for the CCDF. There is little evidence in the literature on selection by initial child skill endowments into CCDF application and receipt and the model baseline exhibits only slight positive selection along this margin. This selection is driven by higher investment in children with higher initial skill endowments, a mechanism consistent with the findings of Datar et al. (2010), Frijters et al. (2013), Grätz and Torche (2016), and Attanasio et al. (2020), among others.

5.2. Distributions of CTC and CCDF spending and recipients

Fig. 7 shows how spending on the CTC and CCDF is allocated across pretax income quartiles for families with children under 5 in the baseline equilibrium, which is the population of interest for child skill accumulation in our model environment. It also includes empirical counterparts for these spending patterns, replicated from Fig. 1 of Section 2. These spending patterns are useful to determine whether the model baseline captures the relative degree of means-testing of the CTC and CCDF, given our parameterization choices, the endogenous income distribution, and the endogenous decision to apply for and take up the CCDF subsidy and how much childcare to use upon receipt. As for the remaining two family policies, Appendix C.2.4 shows the distribution of spending for the

[12] The patterns of selection into CCDF receipt documented here are also apparent if we use realized income to define eligibility.

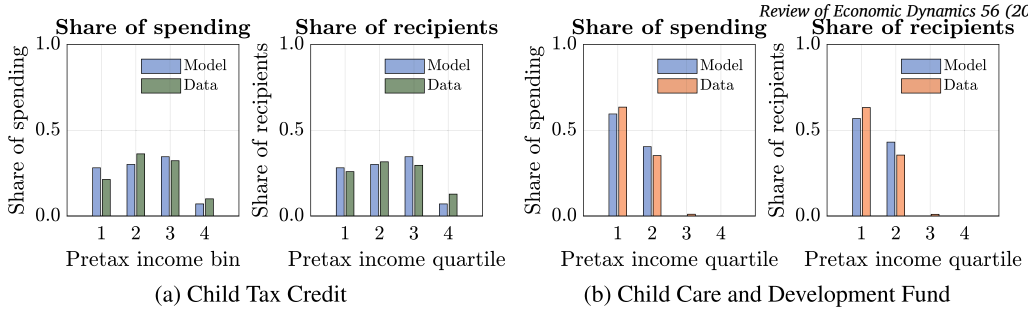

Fig. 7. Distributions of spending and recipients by income quartile. **Notes:** Fig. 7 presents, for households with children under 5 in the model baseline equilibrium, the share of program spending (left panel) and share of recipients (right panel) going to each income quartile for the CTC (Fig. 7a) and the CCDF (Fig. 7b). The y-axis of each graph is the share of total spending or share of total recipients; the x-axis is the income quartile assigned using the distribution of income for households with children under 5.

Table 3
Government spending in model baseline and in US data.

Policy	Data	Model	Data source
Child Tax Credit (CTC)	0.27	0.27	IRS SOI (2017), Table 3.3
Child Care and Development Fund (CCDF)	0.03	0.06	Chien (2019a,b, 2020)
Earned Income Tax Credit (EITC)	0.36	0.39	IRS SOI (2017), Table 3.3
Temporary Assistance for Needy Families (TANF)	0.02	0.02	Office of Family Assistance (2015)

Notes: Table 3 reports the spending on each family policy in the baseline economy and in US data during the period 2015-2017, as well as the source of this empirical moment. All spending amounts are expressed as a percentage of output (model) or US GDP (data).

EITC, which aligns well with the data. We do not report the spending distribution of TANF because it is concentrated in the bottom income quartile by construction (only non-working parents with children may receive TANF in our model).

Fig. 7a illustrates the distribution of spending on the CTC. The model captures the property that the second and third income quartiles together receive most of CTC spending, both in dollar terms (left panel) and as a share of recipients (right panel). Fig. 7b illustrates the distribution of spending on the CCDF. The model captures the fact that spending is concentrated in the bottom two income quartiles for this policy, both in dollar terms (left panel) and in terms of the distribution of recipients (right panel). Compared to the CTC, in both the model and the data the CCDF is more means-tested in the sense that spending goes to relatively poorer parents with young children.[13]

5.3. Government spending on family policies

Table 3 reports the size of government spending on the four family policies (CTC, CCDF, EITC, and TANF) in the model's baseline equilibrium and in the data. This comparison is useful because it verifies that the size of these policies relative to the rest of the economy is reasonable. Empirical moments are computed as explained in Appendix A.1; all spending amounts are expressed as a percentage of output (model) or of average US GDP from 2015 to 2017 (data) computed from BEA Table 1.1.5 (2022a). Note that spending on SS pensions, the remaining transfer policy, is targeted in our model calibration.

Although the moments of Table 3 are untargeted, the model aligns rather well with the data. The model slightly overstates spending on the EITC and CCDF; proportionally, the largest difference is in total CCDF spending. To investigate the source of this discrepancy, in Table C.13 of Appendix C.2.2, we compare CCDF recipient attributes in the model baseline with observable moments computed using CCDF administrative data combined with information from CPS ASEC. We find that the income of the average CCDF recipient, compared to the income of the average parent with a child under 5, is similar in the model baseline and in the data. Conditional on receiving the subsidy, the magnitude of the transfer is slightly higher in the model than in the data; this is partly due to the subsidy intensity (which we approximate with a linear function) being slightly higher conditional on receipt than in the data, but it is likely also due to our abstracting from heterogeneity in childcare prices which tend to be lower among the poor (Moschini, 2023). As for the mass of transfer recipients for each family policy, a comparison with the data is discussed in Appendix C.2.3. For untargeted moments where it is possible to compute the empirical counterpart (namely, the CTC and EITC), the model baseline exhibits a good fit to the data along this dimension.

[13] To explore an alternative model specification with different implications for sorting into CCDF receipt and the distribution of aggregate spending, in Appendix C.5.3 we posit a model where application costs are in utils, recalibrate the model, and demonstrate that selection favors lower-skilled adults in the baseline equilibrium (in contrast to our preferred specification with application costs in the budget constraint). With application costs in utils, the CCDF spending distribution is counterfactually concentrated in the first income quartile. We repeat our main experiments in the modified framework; skill gains from the CCDF expansion are higher, but our main results are qualitatively unchanged.

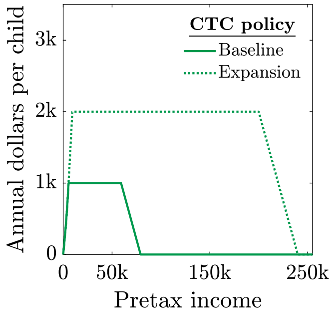

Fig. 8. CTC transfer level in the baseline and after its expansion. **Notes:** Fig. 8 illustrates the CTC in the baseline (solid line) and after its expansion (dotted line). Income and transfer dollar values are expressed in 2016 US dollars by scaling using median household income.

6. Experiments and results

This section compares the effects of spending-equivalent expansions of the CTC and the CCDF. Section 6.1 describes how the policy expansions are implemented; Section 6.2 explains our three main measures of welfare changes. In Section 6.3.1, we report changes across stationary steady states in aggregate variables and welfare, while in Section 6.3.2 we discuss changes in the distribution of adult skill and sources of gains in average adult skill. Section 6.3.3 decomposes welfare gains behind the veil of ignorance from each expansion into marginal contributions from four components.

6.1. Implementation of policy expansions

The CTC expansion is based on the 2017 Tax Cuts and Jobs Act changes to the CTC, described in Crandall-Hollick (2018): the Act increased both the maximum nonrefundable and refundable credit levels, increased the phase-out threshold, and decreased the refundability threshold, which we convert into changes in $\Theta_{C,\kappa,n}$, $\Theta_{C,\kappa,r}$, $\Theta_{C,y,n}$, and $\Theta_{C,y,r}$.[14] Fig. 8 illustrates how this expansion affects our model representation of the CTC, scaled into US dollars using median household income. The CCDF expansion is implemented by eliminating rationing (that is, setting $\phi_N = 1$ so that application guarantees receipt) and subsidizing application costs by raising the policy instrument κ_N from equation (10) until total government spending on all expenditures, expressed as a percent of GDP, is equivalent to the CTC expansion. This equivalence in government spending is how we render the two expansions comparable.

6.2. Measures of welfare changes

We measure welfare changes in consumption-equivalent units in three ways: (a) behind the veil of ignorance; (b) on average for new adults (age $j = 1$) who know their own skill; and, (c) on average across all adults who know their own skill. For (b) and (c), we also report the share who experience welfare gains; no one is indifferent. Here, we extend the notation of the model section to reflect the steady state equilibrium associated with a given government policy: Φ denotes the baseline equilibrium and Φ' denotes the new equilibrium after the policy change. Expressions for consumption-equivalent welfare changes reflect the utility function of Section 3.

Behind the veil of ignorance Welfare changes behind the veil of ignorance are constructed in the following way. First, we construct expected utility behind the veil of ignorance in the baseline equilibrium by integrating the expected utility at stage 1 of period $j = 1$, $V_1\left(\theta_a \mid \Phi\right)$ as defined in equation (1), over θ_a using the distribution of adult skill in the baseline equilibrium: $\mathbb{E}_{\theta_a} V_1\left(\theta_a \mid \Phi\right) = \int_{\theta_a} V_1\left(\theta_a \mid \Phi\right) \mu\left(\theta_a \mid \Phi\right) d\theta_a$. Similarly, the expected utility behind the veil of ignorance in the new steady state is: $\mathbb{E}_{\theta_a} V_1\left(\theta_a \mid \Phi'\right) = \int_{\theta_a} V_1\left(\theta_a \mid \Phi'\right) \mu\left(\theta_a \mid \Phi'\right) d\theta_a$. We then calculate the consumption-equivalent change in the expected lifetime utility behind the veil of ignorance that results when moving from the baseline equilibrium Φ to the new equilibrium Φ':

$$W_V\left(\Phi, \Phi'\right) = 100 \times \left[\exp\left(\frac{\mathbb{E}_{\theta_a} V_1\left(\theta_a \mid \Phi'\right) - \mathbb{E}_{\theta_a} V_1\left(\theta_a \mid \Phi\right)}{\sum_{age=1}^{J} \beta^{age-1}}\right) - 1\right] \tag{15}$$

When we report welfare changes "behind the veil of ignorance" (measure a) we report $W_V\left(\Phi, \Phi'\right)$ directly. This welfare measure represents changes in the expected utility of being a new adult entering the labor market in the final equilibrium as opposed to the initial equilibrium, with no information about one's own skill level: $W_V\left(\Phi, \Phi'\right)$ does not take individual attributes as arguments because it does not condition on them.

[14] The specifics of this conversion are as follows. The changes to the CTC affected both single tax filers and those married filing jointly. As with the baseline CTC parameterization, we average across tax filing statuses using population weights, so that in dollar terms the changes are such that: the expansion increases the maximum credit level, $\overline{\kappa}_n$, from \$1,000 to \$2,000; increases the maximum refundable credit level, $\overline{\kappa}_r$, from \$1,000 to \$1,400; decreases the refundability threshold, $\underline{y}_{C,r}$, from \$1,815 to \$1,513; and increases the phase-out threshold, $\underline{y}_{C,n}$, from \$59,200 to \$200,000. We convert dollar changes to changes in $\Theta_{C,\kappa,n}$, $\Theta_{C,\kappa,r}$, $\Theta_{C,y,n}$, and $\Theta_{C,y,r}$ by dividing by median household income, estimated as described in Section 4.

Adults with a given skill level Welfare changes for adults at each age are constructed holding fixed age j and adult skill θ_a. When computing these changes, for age $j = 1$ we use the stage 1 expected value function of a new adult defined in equation (1), which only depends on θ_a. For older ages, there is no risk and we use the value function for that age which is indexed by only age and adult skill. Specifically, this welfare measure is constructed as follows:

$$W_A\left(j,\theta_a,\Phi,\Phi'\right) = \begin{cases} 100\times\left[\exp\left(\frac{V_1\left(\theta_a|\Phi'\right)-V_1\left(\theta_a|\Phi\right)}{\sum_{age=j}^{J}\beta^{age-j}}\right)-1\right] & \text{if } j=1 \\ 100\times\left[\exp\left(\frac{V\left(j,\theta_a|\Phi'\right)-V\left(j,\theta_a|\Phi\right)}{\sum_{age=j}^{J}\beta^{age-j}}\right)-1\right] & \text{otherwise} \end{cases} \tag{16}$$

When we report welfare changes for "New adults," (measure b) we refer to the average of (16) across θ_a when $j = 1$, $\int_{\theta a} W_A\left(1,\theta_a,\Phi,\Phi'\right)\mu\left(\theta_a \mid \Phi'\right) d\theta_a$, where we integrate using the distribution over adult skill levels in the new equilibrium. The share of new adults who experience welfare gains integrates over an indicator for positive welfare changes: $\int_{\theta a} \mathbb{I}_{W_A(1,\theta_a,\Phi,\Phi')>0}\mu\left(\theta_a \mid \Phi'\right) d\theta_a$.

By welfare changes for "All adults," (measure c) we refer to $\sum_{j=1}^{J}\frac{1}{11}\int_{\theta a} W_A\left(j,\theta_a,\Phi,\Phi'\right)\mu\left(\theta_a \mid \Phi'\right) d\theta_a$, which is computed by integrating (16) for each age using the distribution over adult skill levels in the new equilibrium and then averaging across ages. The share of all adults who experience welfare gains is: $\sum_{j=1}^{J}\frac{1}{11}\int_{\theta a} \mathbb{I}_{W_A(j,\theta_a,\Phi,\Phi')>0}\mu\left(\theta_a \mid \Phi'\right) d\theta_a$.

Note that, in contrast to the welfare function computed behind the veil of ignorance in equation (15), the welfare measure of equation (16) represents changes in well-being for adults who internalize their skill level but take expectations over other elements of their type space. For $j = 1$, the welfare measure in equation (16) incorporates the fact that, once an adult knows their skill level, they do not value a higher average skill among adults as they would behind the veil of ignorance; instead, a higher average skill level affects them indirectly, for example through the income tax rate. In addition, a high-skill adult values expanding transfers to the poor to the extent that it affects their potentially low-skill children, but understands that they will not benefit from it directly (and may have to pay higher taxes). For older ages, the welfare measure of equation (16) also incorporates losses to adults who are too old to benefit from expanded family policies but nevertheless have to pay a higher tax rate compared to the initial equilibrium, or face a different pension level due to changes in labor supply. From a political economy standpoint, of the welfare change metrics that we use in this section, we consider the share of all adults who benefit from a policy change to most closely reflect the preferences of voters in the economy in a given equilibrium.

6.3. Results

6.3.1. Changes in aggregate variables and welfare

Table 4 reports changes across steady states in aggregate variables and welfare. First, we report levels of each variable in the baseline equilibrium. Next, column (1) contains changes after the CTC expansion; in column (2), we report changes after the expansion in the CCDF. The units of these changes are noted in each panel: Panel A reports percentage point changes in government spending as a percent of output and the income tax rate evaluated at the baseline equilibrium's average pretax income level; Panel B reports percent changes in aggregate quantities, specifically average adult skill, average labor supply, output, and average consumption; using the three welfare measures explained in Section 6.2, Panel C reports welfare changes, and Panel D reports the share of new adults and all adults who gain in each equilibrium (no one is indifferent, so the share who lose is one minus this share).

Table 4
Changes in aggregates and welfare across steady states.

Variable category	Variable description	Baseline	(1) CTC	(2) CCDF
Panel A: Government	Government expenditures	22.62	0.54	0.54
Δ units: percentage point change	Average tax rate (at y_{ave} in baseline)	17.68	0.53	0.23
Panel B: Aggregate quantities	Average adult skill	0.94	0.51	6.85
Δ units: percent change	Labor supply	0.31	0.14	-0.61
	Output	0.43	0.52	6.61
	Consumption	0.35	0.49	6.02
Panel C: Welfare change in levels	Behind the veil of ignorance	0.00	2.45	8.28
Δ units: percent of lifetime consumption	New adults	0.00	1.87	2.11
	All adults	0.00	-0.18	0.00
Panel D: Share who gain	New adults	0.00	1.00	0.66
Δ units: share of group	All adults	0.00	0.22	0.15

Notes: Table 4 reports baseline values and changes in aggregates and welfare predicted by the model after the expansion of the CTC (column 1) and a spending-equivalent expansion of the CCDF (column 2). Specifically, in column (1) changes are as described in the main text, and in column (2) $\kappa_N = 0.375 \times \mu_N$ and $\phi_N = 1$.

Beginning with Panel A of column (1), the CTC expansion increases total government spending by 0.54 percentage points; total government spending includes spending on all transfer programs and government consumption, and reflects changes in any of those expenditures that may result from the policy change. To balance the government's budget, the income tax increases; evaluated at the average income in the baseline equilibrium, this increase is 0.53 percentage points. In Panel B, the average skill of adults increases by 0.51 percent, because some of the additional transfers received by parenting families are directed toward investment in child skill. Average labor supply rises very slightly, by 0.14 percent: this net change reflects an increased incentive to move into work for families whose children are young enough (due to the increased generosity of the CTC), and an incentive to distort labor supply downwards to qualify for a higher transfer from the credit (which is hump-shaped in income as shown in Fig. 8) as well as in response to higher income taxes. Overall, output increases by 0.52 percent and consumption increases by 0.49 percent. Welfare gains in Panel C are 2.45 percent behind the veil of ignorance, 1.87 percent on average across new adults, and -0.18 percent on average across all adults. Gains behind the veil of ignorance are higher than gains for new adults because the former internalizes increases in average adult skill, while gains for new adults are higher than gains for all adults because older consumers pay higher taxes without receiving transfers and their welfare declines in equilibrium.[15] Panel D reflects redistribution across ages: all new adults benefit from the CTC expansion, but only about 1 in 5 of all adults in the economy do (recall that, in the model, the CTC is available to adults until their child turns 15, contingent on meeting income requirements).

Moving to Panel A of column (2), the CCDF expansion increases total government spending by the same amount as in column (1) by setting $\phi_N = 1$ and $\kappa_N = 0.375 \times \mu_N$. However, the government's budget is balanced with a tax rate increase of only 0.23 percentage points, lower than the tax increase in column (1). This is possible because the increase in the average skill of adults is higher, at 6.85 percent (Panel B). Average labor supply declines slightly, by 0.61 percent; this change is the net effect of some adults distorting their labor supply downward in order to qualify for the CCDF and other adults raising labor supply by joining the labor force to meet the CCDF work requirement, as well as downward adjustments in response to increased tax rates. Overall, output increases by 6.61 percent and consumption by 6.02 percent. Welfare gains in Panel C are 8.28 percent behind the veil of ignorance, 2.11 percent averaging across new adults, and about zero averaging across all adults. Gains behind the veil of ignorance are higher than in column (1) because of the larger increase in the average skill of adults; average welfare gains for new adults and all adults are similar across the two expansions, although slightly higher for the CCDF. In Panel D, a smaller share of new adults and all adults benefit from the CCDF expansion (66 and 15 percent, respectively) compared to the CTC expansion of column (1). Note that the increase in the size of the economy in column (2) means that other transfer programs also expand, which acts to boost the share of adults who experience welfare gains in column (2).

To summarize, behind the veil of ignorance, the CCDF expansion is preferred because of larger increases in average adult skill. New adults prefer the CTC expansion because a larger share of new adults experience welfare gain, although on average the level of gains is slightly lower than gains from the CCDF. The majority of adults experience welfare losses from expansions in family policies, because in our model most adults do not have children living with them anymore and therefore do not receive transfers from family policies and do not anticipate receiving such transfers in the future; at the same time, they must pay higher taxes. From a political economy standpoint, if the adult population in the final steady state voted, neither the CTC nor the CCDF would be expanded. If adults had to choose to expand one policy or the other, they would choose the CTC.

Our method for expanding each policy utilizes more than one policy margin. In Appendix C.3.2 we examine the role of each margin: for the CTC expansion, the rise in the credit level and phaseout threshold drive most of the welfare gains, while for the CCDF expansion the application cost subsidy makes the largest contribution. Here, we have also targeted a specific change in government spending; in Appendix C.4.1, we gradually expand each policy to show that the pattern of welfare gains documented in Panel C of Table 4 holds true over a wide range of spending-equivalent expansions.

As for the specific policy margins of each expansion, the changes to the CTC are drawn from legislation, while the combination of policy margins used to expand the CCDF is our modeling choice. In Appendix C.4.2, we compare the CCDF expansion of the main text with two alternative spending-equivalent CCDF expansions: first, one that only subsidizes application costs; and, second, an expansion that maintains work requirements but eliminates means testing by removing the income eligibility threshold and setting the subsidy to a flat rate which is calibrated to target government spending. The second alternative expansion represents a more traditional expansion of the CCDF, in the sense that it does not reduce the frictions associated with implementation via a spending program.[16] Of these three CCDF expansions, that of column (2) in Table 4 yields the largest welfare gains and the traditional expansion yields the lowest gains. This is because reducing application costs generates CCDF inflows from lower-skill adults, while a traditional expansion that leaves these application costs unchanged generates inflows from higher-skilled adults who can afford to pay application costs. Thus, after a traditional expansion, the poor face an increased tax burden that is not accompanied by increased transfers, which reduces welfare gains.

6.3.2. Changes in the distribution of adult skill and sources of skill gains

As shown in Panel A of Table 5, the variance of log skill increases after both expansions, although to a larger extent after the CCDF expansion. This increase indicates that the rise in average adult skill is not shared equally over the skill distribution. To attribute these changes in the variance to parts of the skill distribution, we first divide the distribution of log skill into deciles and then take the difference between the 10th and 1st, 10th and 5th, and 5th and 1st within-decile averages in each equilibrium. Panel B reports the changes in these statistics compared to the baseline: both policies give rise to an increase in inequality because lower percentiles of the skill distribution see small gains relative to higher percentiles. After the CTC expansion, inequality among higher percentiles

[15] Technically, the EITC in our model economy makes small transfers to those who have no children at home. However, this portion of the EITC has stringent eligibility requirements and is received by few adults.

[16] We could instead subsidize the application costs in combination with a different policy margin. To motivate our choice in this respect, in Appendix C.3.3 we compare the welfare gains of expansions of the CCDF before subsidizing application costs; of the margins we consider, eliminating rationing generates the largest welfare gain.

Table 5
Changes in the distribution of adult skill across steady states.

Variable category	Variable description	Baseline	(1) CTC	(2) CCDF
Panel A: Mean and overall inequality	Average	0.938	0.510	6.849
Δ units: percent change (average) and levels (variance of log)	Variance of log	0.654	0.002	0.038
Panel B: Decile dispersion (log skill)	10th-1st	2.604	-0.001	0.032
Δ units: levels	10th-5th	1.308	-0.014	-0.018
	5th-1st	1.296	0.013	0.050

Notes: Table 5 reports statistics related to the distribution of skill in the baseline and changes relative to the baseline in these statistics after an expansion in (1) the CTC and (2) the CCDF.

decreases while inequality among lower percentiles increases by a similar amount. By comparison, the increase in inequality among lower percentiles of the skill distribution after the CCDF expansion is much greater in magnitude, indicating that the lowest skill percentiles are left behind to a greater extent and driving the larger increase in variance of log skill for this policy.

In Appendix C.2.1, we show that the average child skill outcome is increasing in the child's initial skill endowment in the baseline equilibrium. However, the gradient of this increase is less steep for children with low-skill parents, indicating that there is room for improvement in the outcomes of children in high-θ_k, low-θ_a families. In Appendix C.3.4, we show that gains in child skill outcomes under both the CTC and CCDF expansions are concentrated among families with low θ_a and high θ_k types.[17] The mass of these high-skill-gain families is determined by the correlation between θ_k and θ_a, which suggests that this correlation is related to the overall changes in average adult skill after each policy is expanded. In our benchmark calibration, our target value for this correlation is set to 0.06, the correlation between child and mother cognitive skill reported in the online appendix of the well-known study of Cunha et al. (2010); the correlation is computed using nationally-representative US data collected when the child is around 12 months old. We choose this value for our calibration target because it reflects nationally representative data, the child is young enough to be interpreted as close to birth, and because the target value lies between the values of conceptually similar statistics computed in Moschini (2023) (using data on 9-month-old children) and in Attanasio et al. (2020) (using data on 18-month-old children). The estimates from these three studies differ in the data set used, measures of skill, and age of the children when the skills are measured; with such caveats in mind, the overall pattern that emerges is that the measured correlation of child skill and parent skill or income is far less than 1 and is increasing as the child progresses through early childhood.

This pattern is consistent with a low initial correlation combined with endogenous skill investment. One concern, however, is that measurement error at very early ages may bias the estimated correlation downwards; at the same time, a calibration target measured later during early childhood is a conceptually less accurate empirical counterpart to the model object we seek to discipline. Putting the latter concern to the side, a higher initial correlation coefficient would reduce the mass of families that experience high skill gains from the CCDF expansion and possibly narrow the gap in welfare gains behind the veil of ignorance between the CTC and CCDF expansions. To investigate this mechanism, in Appendix C.5.1 we recalibrate to the higher correlation coefficient of Attanasio et al. (2020), equal to 0.18, and rerun our main experiments. As expected, we find that in the recalibrated environment the CCDF expansion yields lower gains in average adult skill and lower welfare gains behind the veil of ignorance, compared to the CCDF expansion of Table 4. However, qualitatively the takeaways are unchanged: compared to the CTC expansion, adult skill gains are higher after the CCDF expansion, the CCDF is preferred to the CTC behind the veil of ignorance, and adults who know their own type on average experience similar welfare gains across the two policies.[18]

6.3.3. Decomposition of welfare changes behind the veil of ignorance

To identify the sources of welfare gains from each policy expansion, we decompose welfare changes behind the veil of ignorance from Panel C in Table 4 into marginal contributions from four components: (i) changes in the age profile of consumption while holding fixed aggregate consumption at its baseline value; (ii) changes in the average altruism term value due to changes in child skill outcomes; (iii) changes in the average levels of consumption and non-leisure time; and, (iv) a distributional component that embeds remaining adjustments, including changes in the effect of skill on expected lifetime utility in the altruism term of parents, how quantities are distributed across types, and changes in the distribution of types in the economy. Here, due to the policies we examine, we isolate the effect of changes in the age profile of consumption (component i), generalizing the welfare decomposition of Guvenen et al. (2023a); the implementation of our decomposition is described in Appendix B.4. Table 6 presents the decomposition results.

For the CTC, the largest contributions to welfare gains stem from changes in the distributional component (component iv) followed by the age profile of consumption (component i). Changes in the age profile of consumption (component i) reflect redistribution of consumption towards younger adults, facilitated by the CTC. The distributional component (component iv) embeds several adjust-

[17] In Appendix C.3.4 we also report gains in child skill outcomes when we additionally eliminate means-testing after the CCDF expansion. Even when high-θ_a families have access to childcare subsidies, high-skill-gain families remain concentrated among high-θ_k, low-θ_a families.

[18] Because measured correlation coefficients tend to increase as the child ages, there may be higher estimates of the correlation between the skills of young children and those of their parents, not discussed in Appendix C.3.4, when measured even later during early childhood. At some point, however, concerns about reasonably interpreting the correlation coefficient as measuring initial conditions start to outweigh concerns about measurement error, especially if one believes (consistent with the premise of this paper) that endogenous investment affects child skill accumulation.

Table 6
Contributions to welfare gains behind the veil of ignorance.

Policy	(i) Age	(ii) Child	(iii) Levels	(iv) Dist.	Total
CTC	0.33	0.06	0.18	0.42	1.00
CCDF	-0.02	0.19	0.84	-0.02	1.00

Notes: Table 6 decomposes the sources of welfare gains behind the veil of ignorance across steady states after the CTC expansion (first row) and the CCDF expansion (second row). Columns (i) to (iv) marginally introduce an additional change in the economy; for the welfare metric in each row, the cell reports the share of the total welfare change explained by the contribution of the column. The marginal contributions are described in the main text. Column titles are abbreviations of what each component represents; see the text and Appendix B.4 for details.

ments, but here mostly reflects changes in the distribution of types and changes in the effect of a given child skill outcome on expected lifetime utility in the altruism term of parents. The former effect is positive because the adult skill distribution shifts towards types that benefit more from the CTC expansion, compared to the baseline distribution. The latter effect is positive because adults internalize that, at a given skill outcome, their child will also benefit from the CTC.[19] Changes in average child skill outcomes and average levels of consumption and non-leisure time (components ii and iii, respectively) also make positive, albeit smaller, contributions to overall welfare gains.[20]

For the CCDF, the largest contribution stems from changes in the average levels of consumption and non-leisure quantities (component iii). Specifically, average consumption in the economy as a whole goes up relative to the baseline because adults have higher skill levels on average in the new economy; at the same time, average labor supply slightly decreases (for example because adults age $j = 1$ distort labor downward to qualify for the CCDF). Changes in the average altruism term due to changes in child skill outcomes (component ii) also contribute positively. In contrast to the CTC, changes in the age profile of consumption (component i) and other adjustments (component iv) act to slightly reduce welfare gains from the CCDF expansion. Component (i) is slightly negative because the share of aggregate consumption going to age $j = 1$ decreases as CCDF recipients substitute away from consumption and towards investment in child skill. Contributions from remaining adjustments (component iv) represent the net effect of several changes. A large share of adults experience welfare losses after the CCDF expansion, and the net negative contribution of this component reflects these unequal gains across consumers; at the same time, the other shifts incorporated in component (iv) act to reduce the magnitude of its negative contribution.

Comparing across policies, the contribution of changes in the age profile of consumption is more important for welfare gains from the CTC, while the contribution from changes in the level of consumption is more important for welfare gains from the CCDF. This highlights that the CTC redistributes consumption to earlier in life and thereby addresses the inability of consumers to borrow against their future income to finance consumption when young. By comparison, the CCDF encourages investment in child skill, raising the average level of consumption in the economy by raising the average skill of adults. We intuitively connect this to the inability of children to motivate more investment in their skill from their parents: the CCDF can be viewed as a subsidy to skill investment that acts to partially address this missing market between generations. Additionally, using welfare changes behind the veil of ignorance from Table 4 to scale up the shares of Table 6 suggests that, in our environment, policies which partially compensate for children's inability to motivate additional skill investment (such as the CCDF) will tend to yield larger welfare gains compared to policies that, to a greater extent, partially compensate for life cycle borrowing constraints in adulthood by redistributing consumption to earlier ages (such as the CTC).

In Appendix C.3.5, we examine the contribution of aggregate equilibrium adjustments to welfare changes in general equilibrium. Without equilibrium adjustments in the adult skill distribution, behind the veil of ignorance the CTC would be preferred to the CCDF; this finding is consistent with the marginal contribution of changes in average levels of consumption and non-leisure time documented in the decomposition of Table 6. As for average welfare for new adults and all adults, without equilibrium adjustments in income taxes or the scaling of other transfer policies, the CTC would be preferred (for new adults) or generate almost identical welfare changes (for all adults) compared to the CCDF expansion.

In Appendix C.3.6, we report general-equilibrium changes in our three welfare metrics over the transition to the new steady state. Early in the transition, gains from the CTC expansion exceed those of the CCDF expansion. It takes at least a generation (20 years) after the period in which the policy change is introduced for the CCDF welfare gains to exceed gains from the CTC; the first cohort of treated children with higher skill outcomes need to age into the labor market, have their own children, and start paying taxes. By contrast, gains from the CTC are close to their long-run value soon after the policy change; they are less dependent on the skill composition of adults, which is a slow-moving object. The share of adults (both new adults and overall) who benefit from each policy expansion is higher for the CTC compared to the CCDF, both early in the transition and at the long-run steady state.

[19] As its definition indicates, component (iv) also reflects changes in the allocation of quantities across types, for example shifts to more or less beneficial combinations of consumption and labor supply than what is implied by the proportional changes implemented in other components of the decomposition. However, the contribution of these adjustments to welfare changes is small for the CTC expansion.

[20] Parental altruism means that parents partially internalize the expected lifetime utility of their child at the child skill outcome they choose. This object can increase because the child's skill outcome increases (component ii) or because expected lifetime utility goes up at a given child skill outcome, as reflected by component (iv).

7. Conclusion and directions for future research

In this paper, we compare spending-equivalent expansions of the CTC and CCDF, two major family policies in the United States. Across steady states, the CCDF expansion yields larger welfare gains behind the veil of ignorance compared to the CTC, driven by larger increases in average adult skill. Adult skill inequality increases under both policies, but more so for the CCDF. Average welfare gains for new adults and all adults are similar across the two policy expansions, while a larger share of both new adults and all adults benefit from the expansion of the CTC. To articulate differences in the sources of welfare gains across the two policies, we decompose welfare gains behind the veil of ignorance and find that redistribution of consumption across ages drives a relatively larger share of welfare gains after the CTC expansion, while changes in the average level of consumption and non-leisure time (reflecting higher average adult skill) contribute relatively more to welfare gains from the CCDF expansion.

There is clearly potential for future research on family policies to offer useful insights for policy makers. In our analysis, we have emphasized frictions associated with the implementation method of childcare subsidies in the United States, and explored alternative expansions of family policies taking these frictions into account. In the process, we make several assumptions in our model framework, which include no borrowing or saving, no financial transfers to children, no earnings risk, and a fixed pretax return to skill (wage rate) on the labor market. Future work might consider in more detail the impact of accounting for these features. As suggested by our decomposition of welfare gains from the CTC, allowing consumers to borrow would most likely reduce welfare gains from expanding the CTC by facilitating consumption-smoothing across ages via the private market (although if borrowing on the private market carries a default premium, then the CTC could still offer a less costly way to borrow against one's future income). Allowing consumers to save could also reduce the benefits of family policies: adults might set aside part of the transfer to finance tax liabilities later in life or, if allowed to delay fertility, could accrue savings before parenthood and reduce the need for additional resources when the child is young. Financial transfers from parents to children could be financed with set-aside earnings or government transfers, and could be used to encourage college enrollment (complementing skill investment) or to compensate for low initial skill endowments (substituting skill investment). Earnings risk would relax the link between adult skill and expected lifetime utility and possibly reduce the impact of transfer policies that seek to encourage investment in child skill. Lastly, if the pretax return to adult skill on the labor market is allowed to adjust in equilibrium, this would likely reduce the pass-through of increases in average adult skill to the income tax base, raise the tax burden on older consumers, and lower welfare gains under either family policy expansion that we consider here.

Appendix A. Supplementary material

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.red.2025.101270.

Data availability

Data or data download instructions are included in the replication package associated with this paper.

References

Administration for Children & Families (accessed: 07.15.2024a). FY 2015 Final Data Table 9.

Administration for Children & Families (accessed: 07.15.2024b). FY 2016 Final Data Table 9.

Administration for Children & Families (accessed: 07.15.2024c). FY 2017 Final Data Table 9.

Chien, N., 2019a. Factsheet: Estimates of Child Care Eligibility & Receipt for Fiscal Year 2015. Technical report. Department of Health and Human Services of the United States, Office of the Assistant Secretary for Planning & Evaluation.

Chien, N., 2019b. Factsheet: Estimates of Child Care Eligibility & Receipt for Fiscal Year 2016. Technical report. Department of Health and Human Services of the United States, Office of the Assistant Secretary for Planning & Evaluation.

Chien, N., 2020. Factsheet: Estimates of Child Care Eligibility & Receipt for Fiscal Year 2017. Technical report. Department of Health and Human Services of the United States, Office of the Assistant Secretary for Planning & Evaluation.

Attanasio, O., Cattan, S., Fitzsimons, E., Meghir, C., Rubio-Codina, M., 2020. Estimating the production function for human capital: results from a randomized controlled trial in Colombia. Am. Econ. Rev. 110 (1).

Bastani, S., Blomquist, S., Micheletto, L., 2020. Child care subsidies, quality, and optimal income taxation. Am. Econ. J.: Econ. Policy 12 (4), 1–37.

Bureau of Labor Statistics, US Department of Labor, 2024. Consumer Price Index for All Urban Consumers: All items in US city average, urban consumers, not seasonally adjusted, annual average [CUUR0000SA0]. (accessed 05.23.2024).

Chetty, R., Hendren, N., Kline, P., Saez, E., 2014. Where is the land of opportunity? The geography of intergenerational mobility in the United States. Q. J. Econ. 129 (4), 1553–1623.

Crandall-Hollick, M.L., 2018. The Child Tax Credit: Current Law. Congressional Research Service of the United States R41873.

Crandall-Hollick, M.L., 2021. The Child Tax Credit: Temporary Expansion for 2021 Under the American Rescue Plan Act of 2021 (ARPA; P.L. 117-2). Congressional Research Service of the United States IN11613.

Crandall-Hollick, M.L., Falk, G., Boyle, C.F., 2021. The Earned Income Tax Credit (EITC): How It Works and Who Receives It. Congressional Research Service of the United States R43805.

Cunha, F., Heckman, J., Schennach, S., 2010. Estimating the technology of cognitive and noncognitive skill formation. Econometrica 78 (3), 883–931.

Daruich, D., 2023. The Macroeconomic Consequences of Early Childhood Development Policies. Working Paper.

Datar, A., Kilburn, M.R., Loughran, D.S., 2010. Endowments and parental investments in infancy and early childhood. Demography 47 (1), 145–162.

Flood, S., King, M., Rodgers, R., Ruggles, S., Warren, J.R., Backman, D., Chen, A., Cooper, G., Richards, S., Schouweiler, M., Westberry, M., 2024. IPUMS CPS: Version 11.0 [dataset] (accessed: 01.06.2024).

Frijters, P., Johnston, D.W., Shah, M., Shields, M.A., 2013. Intrahousehold resource allocation: do parents reduce or reinforce child ability gaps? Demography 50 (6), 2187–2208.

Goldin, J., Michelmore, K., 2022. Who benefits from the child tax credit? Natl. Tax J. 75 (2), 123–147.

Grätz, M., Torche, F., 2016. Compensation or reinforcement? The stratification of parental responses to children's early ability. Demography 53 (6), 1883–1904.

Guner, N., Kaygusuz, R., Ventura, G., 2020. Child-related transfers, household labour supply, and welfare. Rev. Econ. Stud. 87, 2290–2321.

Guvenen, F., Kambourov, G., Kuruscu, B., Ocampo, S., Chen, D., 2023a. Use it or lose it: efficiency and redistributional effects of wealth taxation. Q. J. Econ. 138 (2), 835–894.

Guvenen, Fatih and Pistaferri, Luigi and Violante, Gianluca (accessed: 05.15.2023). Global Repository of Income Dynamics (GRID) Database.

Guzman, G.G., 2017. Household Income: 2016. American Community Survey Briefs. ACSBR/16-02.

Guzman, G.G., 2018. Household Income: 2017. American Community Survey Briefs. ACSBR/17-01.

Guzman, J.J., 2019. The demand for child care subsidies under rationing. Rev. Econ. Househ. 17, 1349–1379.

Hagenaars, A.J., de Vos, K., Zaidi, M.A., 1994. Poverty statistics in the late 1980s: research based on micro-data. Stat. Rep. EUROSTAT.

Heathcote, J., Storesletten, K., Violante, G.L., 2017. Optimal tax progressivity: an analytical framework. Q. J. Econ. 132, 1693–1754.

Ho, C., Pavoni, N., 2020. Efficient child care subsidies. Am. Econ. Rev. 110 (1), 162–199.

Huggett, M., Ventura, G., Yaron, A., 2006. Human capital and earnings distribution dynamics. J. Monet. Econ. 53 (2), 265–290.

Internal Revenue Service of the United States, Statistics of Income, 2016. Table 3.3 - All Returns: Tax Liability, Tax Credits, and Tax Payments (2016, accessed: 05.15.2023a).

Internal Revenue Service of the United States, Statistics of Income, 2017. Table 3.3 - All Returns: Tax Liability, Tax Credits, and Tax Payments (2017, accessed: 05.15.2023b).

Johnson, A.D., Martin, A., Brooks-Gunn, J., 2011. Who uses child care subsidies? Comparing recipients to eligible non-recipients on family background characteristics and child care preferences. Child. Youth Serv. Rev. 33 (7), 1072–1083.

Kamerman, S.B., 2000. Early childhood education and care: an overview of developments in the OECD countries. Int. J. Educ. Res. 33, 7–29.

Keane, M., Rogerson, R., 2015. Reconciling micro and macro labor supply elasticities: a structural perspective. Annu. Rev. Econ. 7, 89–117.

Lagakos, D., Moll, B., Porzio, T., Qian, N., Schoellman, T., 2018. Life cycle wage growth across countries. J. Polit. Econ. 126 (2).

Leibowitz, A., 1974. Home investments in children. J. Polit. Econ. 82 (2), 111–131.

Lynch, K.E., 2022. Child Care Entitlement to States: an Overview. Congressional Research Service of the United States IF10511.

Marr, C., Huang, C.-C., Sherman, A., DeBot, B., 2015. EITC and Child Tax Credit Promote Work Reduce Poverty, and Support Children's Development, Research Finds. Center on Budget and Policy Priorities Report.

Mendoza, E.G., Razin, A., Tesar, L.L., 1994. Effective tax rates in macroeconomics: cross-country estimates of tax rates on factor incomes and consumption. J. Monet. Econ. 34, 297–323.

Moschini, E.G., 2023. Child care subsidies and child skill accumulation in one- and two-parent families. Am. Econ. J. Macroecon. 15 (1), 475–516.

Office of Family Assistance, 2015. Temporary Assistance for Needy Families: 12th Report to Congress Fiscal Years 2014 and 2015. Technical report, Department of Health and Human Services of the United States, Administration for Children and Families.

Olivetti, C., Petrongolo, B., 2017. The economic consequences of family policies: lessons from a century of legislation in high-income countries. J. Econ. Perspect. 31 (1).

Social Security Administration of the United States, 2016a. Social Security and Equivalent Railroad Retirement Benefits (SSA Publication 915). Technical report, Social Security Administration of the United States.

Social Security Administration of the United States, 2016b. Your Retirement Benefit: How It's Figured (SSA Publication No. 05-10070). Technical report, Social Security Administration of the United States. January.

Sommartino, F., Toder, E., Maag, E., 2002. Providing Federal Assistance for Low-Income Families Through the Tax System: a Primer. Urban-Brookings Tax Policy Center Working Paper No. 4.

Steuerle, E.C., 1990. Policy watch: tax credits for low-income workers with children. J. Econ. Perspect., 201–2012.

United States Department of Health and Human Services, Administration for Children and Families, Child Care Bureau, 2022. Child Care and Development Fund Administrative Data, Federal Fiscal Year 2016. Inter-university Consortium for Political and Social Research [distributor]. https://doi.org/10.3886/ICPSR37264.v3.

United States Bureau of Economic Analysis (BEA), 2015-2017. Table 1.1.10 Percentage Shares of Gross Domestic Product (accessed: 04.22.2024).

United States Bureau of Economic Analysis (BEA), 2015-2017. Table 2.1 Personal Income and Its Disposition (accessed: 05.15.2023a).

United States Bureau of Economic Analysis (BEA), 2015-2017. Table 1.1.5 Gross Domestic Product (accessed: 05.15.2023b).

Zhou, A., 2022. The Macroeconomic Consequences of Family Policies. Working paper.