BENCHMARKING A CHILD ALLOWANCE AGAINST FREE CHILDCARE

John Walker & Max Ghenis

June 2021

A universal child allowance is considered a simple and effective policy to reduce adult poverty, child poverty, and potentially inequality. It is often compared to childcare subsidies. This study investigates the relative efficacy of provision of free childcare (using a transfer payment equal to the household's expenditure on childcare as a proxy) and of a child allowance with the same fiscal burden (a payment made to each child's guardian equal in value to the value of the total program divided by the number of children). We use US Current Population Survey data aggregated over the 2018-2020 period to simulate the impact of each policy on measures of poverty and inequality, stratifying results by state. We perform a secondary analysis using state-based cost of childcare estimates from the Center for American Progress to represent necessary expenditure on childcare. We find modest improvements in measures of adult poverty and inequality that are substantially larger for the child allowance. We find relatively large reductions in the child poverty rate that are much larger for the child allowance policy. We find substantial heterogeneous effects by state, though the impact of the child allowance is almost unilaterally greater.

Corresponding author: John Walker - jhwalker@mit.edu

Thanks to Ben Boehlert of the UBI Center for providing research assistance.

Contents

1 Motivation

With the advent of the COVID-19 crisis and widespread school closures, there has been increased interest in the potential for provision of child allowances or childcare to offset household consumption losses. The American Rescue Plan, for example, expanded the Child Tax Credit for one year and created a monthly prepayment program slated to roll out in July 2021. President Joe Biden's American Families Plan proposal includes both extending this Child Tax Credit expansion as well as a number of childcare subsidies. Republicans generally prefer cash assistance to childcare subsidies; for example, Senator Mitt Romney's Family Security Act would replace a number of tax deductions and welfare programs with a single expanded child allowance.

How should we compare these policies targeted at assisting families through childcare? Given they are largely intended to assist struggling households, one candidate criterion is their impact on measures of poverty and inequality. This study aims to benchmark the provision of childcare (a transfer equal to the value of expenditure on childcare serves as a proxy) against a child allowance of equal value on these grounds.

2 Background

In the United States, there are multiple programs that assist guardians in paying for childcare. The most expensive is the Child and Dependent Care Tax Credit, which is a 50% tax credit for those making under $125,000, phasing out to a 20% credit at $183,000 remaining at a 20% credit until $400,000 in income, at which point it fully phases out at $450,000. In order to receive the credit, all caregivers must be working. Before 2021, the credit was non-refundable, but the American Rescue Plan made it fully refundable for the year, and The American Families Plan intends to make it permanently refundable. Middle income earners benefit most from the credit as lower earners tend to have lower childcare expenses (providing childcare themselves) and, while high earners expend the most on childcare, benefits phase out with income.

The dependent care flexible spending account (FSA) also allows guardians to save up to $10,500 every year in accounts that are not be subject to income or payroll taxes.

3 Data

Our analysis uses two distinct datasets: the "CPS" and "CAP" datasets. The CPS dataset is the Current Population Survey March Supplement covering the period 2018-2020, retrieved through IPUMS. CPS data contains observations at both the individual and household level.[1] Household-level variables include total resources, which measures the sum of cash income and in-kind benefits net the value of tax liability, work, childcare, medical expenses and child support payments; and the poverty threshold, which is set with reference to the number of people in the family,

[1] We use "household" to refer to Supplemental Poverty Measure unit, which is a group of individuals who live together and are believed to share resources, for the purposes of calculating the Supplemental Poverty Measure.

geographic differences in cost of living and ownership of assets. A person is in poverty if their household's poverty threshold exceeds its resources. We disaggregate household-level indicators to the individual level for inference.

The "CAP" dataset is provided by the Center for American Progress and estimates average childcare costs for 2018. It separates out the expected monthly cost of "base quality" and "high quality" childcare depending on child age and state of residence. Base quality childcare meets minimum state licensing standards, whereas high quality childcare meets stricter standards based on the number of classrooms, teacher-student ratios, teacher salaries and the like. We annualize the estimated childcare costs and merge the CAP and CPS datasets on state and child age.

4 Methodology

In order to analyze the differential distributional effects of reimbursing household childcare expenses (a proxy for provision of childcare) and provision of a per child allowance that only differs by child age, we conduct six simulations using CPS and CAP data. We simulate the household-level transfer under each hypothetical policy, then add it to the existing household resources, allowing for analysis of the impact of the simulated transfers on poverty, deep poverty and inequality. We disaggregate these measures by key demographics of interest. Finally, as an exploratory analysis, we consider potential heterogeneity in program takeup using CPS demographic data.

4.1 Simulating the impacts of reimbursing childcare expenses and a child allowance of equal cumulative value assuming no new takeup using CPS childcare expenditure data

We use CPS data to perform two simulations comparing the distributional impacts of providing transfers to households equal to their expenditure on childcare or instead a child allowance of equal total value. The former case also serves as a proxy for the impact of provision of childcare.

The childcare offset is simply set at the value of existing household childcare expenses with the assumption of no new take up. The flat transfer is allowed to vary by child age, and, to allow it to be flexible to differential takeup by child age, we estimate a simple predictive model of existing childcare expenses for household h depending on the number of children in each age group g.

$$ChildcareExpenditure_h = \sum_{g=1}^{4} \beta_g KidsInAgeCategory_{h,g} + \epsilon_h \tag{1}$$

Where age category g can be infant ($age < 1$), toddler ($1 \leq age < 2$), preschool ($2 \leq age < 5$) and older child ($5 \leq age \leq 12$).

We retrieve estimates of a household's predicted childcare expenses by child age using weighted least squares, aggregating the number of children to the household level and weighting observations by household weights. We then use the β parameter estimates to define the child allowance amounts.

We are interested in preserving comparability between the transfers by having the total program cost be the same between the two simulations $\sum ChildcareExpenditure = \sum ChildAllowance$. However, as we do not include an intercept in our predictive regression (child allowances by definition must be per-child, rather than providing an amount to each household regardless of its number of children), the equivalency statement no longer holds. Hence, we inflate the resulting estimates by a CostRatio - the ratio of the sum of the childcare expenditure cost to the sum of the child allowances as defined above. Finally, we add the value of each transfer to the household-level resources.

4.2 Simulating the impacts of reimbursing estimated childcare expenses and a child allowance of equal cumulative value (at base and high childcare quality) assuming full takeup using merged CPS and CAP datasets

Rather than assuming no new takeup, we turn to the opposite extreme of assuming full take-up of provided childcare or of a child allowance, for comparison. We merge CPS and CAP estimates of childcare costs by child age in a given state[2]. We then generate two sets of simulations. Each set of simulations compares the impact of a transfer equal to the estimated cost of childcare for a child of a given age which differs by state (provided by CAP) to a child allowance. The two sets of simulations differ by assumption on the quality of childcare. The childcare cost by child age in a given state is set at either the "base" or "high" level based on the CAP estimates, as described in the data section. The childcare cost transfer in each simulation set is simply defined as the CAP-estimated cost of childcare for a given child of age a, in state s, of childcare quality q aggregated to the household level h.

$$ChildcareCost_{hsq} = \sum CAPCost_{asq} \tag{2}$$

Note that here h neatly defines the difference in costs between the two sets of simulations. The child allowance is a transfer equal in average value to the average value of the childcare cost variable defined above, i.e. $\sum ChildcareCost_{hsq} = \sum ChildAllowance$. The child allowance flattens variance of the subsidy across states by simply providing an average transfer per age group (at a given quality). We first summate the per child of age category a and quality q childcare costs. We then evenly distribute these age-quality total costs across children of a given age, generating a different simulation for each quality. We then aggregate these costs to the household level to give the household child allowance.

$$ChildAllowance_{hq} = \sum CAPCost_{aq} \tag{3}$$

Note that the definition of the child allowance removes state-based heterogeneity in the transfer amount, allowing for comparison of the distributional impacts of state-induced heterogeneity in childcare costs. Again, the transfers are added to household total resources to allow for analysis of the distributional impacts of the various simulated

[2]We assume no change in childcare costs over the study period (2018-2020)

programs. The simulated scenarios can be summarized as follows:

Table 1: Scenario comparisons

Baseline	**Childcare replacement**	**Budget-equivalent child allowance**	**Datasets**
Baseline	Replace childcare expenses	Child allowance 1 (varies by age)	CPS
Baseline	Base quality childcare	Child allowance 2 (varies by age)	CPS/CAP
Baseline	High quality childcare	Child allowance 3 (varies by age)	CPS/CAP

The associated cost of each scenario[3] is estimated to be:

Table 2: Program cost estimates (USD billions)

Scenario	**Program cost**
Baseline	0
Expense replacement	55
Low quality	229
High quality	426

4.3 Analysing welfare measures

Using the simulated total household resources generated through the processes above, we calculate various measures of poverty using the SPM poverty threshold. We consider both poverty (being below the threshold) and deep poverty (being below half the threshold), disaggregating by state. We also consider the impacts of the simulated policies on child poverty as well as the Gini coefficient, a measure of inequality.

5 Exploratory analysis

5.1 Predictors of nonzero childcare expenditure

We use a linear probability model to estimate the impact of household demographic characteristics on the probability of expending on childcare. We hypothesize that there is substantial heterogeneity in childcare expenditures based on household income and single guardian status, as we might expect that childcare is prohibitively expensive for low income families and single parents households likely tend to have fewer parties among which to share childcare duties.

$$ConsumesChildcare_h = SingleGuardian_h + NumAdults_h + TotalResources_h + \sum_{g=1}^{4} \beta_g \; KidsInAgeCategory_{hg} + \epsilon_h \tag{4}$$

[3] The budget neutral cost with respect to the CAP scenarios was calculated at the person level, whereas for the expense replacement scenario, it was calculated at the household level. Due to different weights, the total transfer sizes differed by roughly 10% and we here report those calculated at the household level for consistency to allow for comparison between simulations. Budget neutrality was preserved within 0.5% both at the person and household level.

Where $ConsumesChildcare$ is a dummy for any household-level expenditure on childcare; $SingleGuardian$ a proxy dummy for whether the household has a single household head; $TotalResources$[4] a measure of household income; and $KidsInAgeCategory_g$ is a vector of child age categories that define how many children of a given age are in a household.

We find that each of the demographics of interest statistically significantly predict whether a household has any annual childcare expenditures. Additional household resources are associated with small (roughly 1.9 percentage points per $1000) increases in the probability of expenditure on childcare which is highly statistically significant ($t \sim 40.6$). Being a single guardian reduces the probability of expending on childcare by about 18 percentage points, and for every additional adult in the household, the probability falls by another 18 percentage points. Households with more than two adults are especially unlikely to purchase childcare. Single and two adult households have equal probability of purchasing childcare all else equal, but it is likely that the number of adults is positively correlated with household resources and other demographic characteristics and therefore that some two guardian households may be more likely to consume childcare.

As expected, having a child of any age group substantially increases the probability of paying any amount for childcare over the study period. The impact of having additional older children or infants in the household on the probability of having childcare expenditures are relatively small, about 7.9 and 10.2 percentage points respectively. This is likely because guardian take on caring duties for the youngest children and older children are more independent and attend school, and therefore have lower childcare needs. By contrast, the probability of paying for childcare given an additional toddler or preschooler is much higher (roughly 18 percentage points), reflecting their greater need for childcare.

[4]Here we use a slightly different measure of household resources, total resources, which does not net out child expenses, to remove endogeneity associated with the other resource metric.

Table 3: Regressing any childcare expenditure on predictive covariates

	Any household childcare expenditures
	(1)
constant	0.038***
	(0.002)
household resources	0.000***
	(0.000)
single guardian	-0.018***
	(0.002)
no. adults	-0.018***
	(0.001)
no. infants ($age < 1$)	0.102***
	(0.003)
no. toddlers ($1 \leq age < 3$)	0.179***
	(0.002)
no. preschoolers ($3 \leq age \leq 5$)	0.176***
	(0.002)
no. older children ($5 < age \leq 12$)	0.079***
	(0.001)
Observations	205,618
R^2	0.196
Adjusted R^2	0.196
Residual Std. Error	5.401(df = 205610)
F Statistic	7143.293*** (df = 7.0; 205610.0)
Note:	*p<0.1; **p<0.05; ***p<0.01

These results provide evidence that, at baseline, use of childcare is a function of household characteristics. It is likely that these characteristics are also somewhat determinative of take up of new childcare opportunities, such as the policies simulated in this analysis. This therefore provides tentative evidence that differential take up will have significant complicating effects on the distributional impact of childcare and child allowance policies.

5.2 Mimicking childcare expenditures with a child allowance

Below, we report the findings of the regression used to predict childcare expenditure by child age. We used the predicted values from this regression to formulate the child allowance for comparison with the replacing childcare expenses simulation.

Table 4: Regressing household childcare expenditures on number of children of a given age group

	Household childcare expenditures
	(1)
no. infants ($age < 1$)	799.891***
	(35.935)
no. toddlers ($1 \leq age < 3$)	1730.534***
	(23.257)
no. preschoolers ($3 \leq age \leq 5$)	1803.074***
	(17.785)
no. older children ($5 < age \leq 12$)	368.430***
	(9.478)
Observations	205,618
R^2	0.126
Adjusted R^2	0.126
Residual Std. Error	63522.545(df = 205614)
F Statistic	7389.380*** (df = 4.0; 205614.0)
Note:	*p<0.1; **p<0.05; ***p<0.01

Annual household childcare expenditures are lowest for older children and also relatively low for infants. Expenditure on toddlers and preschoolers is more than double as much, reflecting high childcare need. These results are consistent with the regression results described in section 4.3.

6 Main results

Here we compare impacts of the three hypothetical childcare policies (reimbursing current expenses, providing funds equal to base-quality childcare in the household's state, and providing funds equal to high-quality childcare in the household's state) with budget-equivalent childcare policies, on poverty and inequality measures.

6.1 Poverty impacts

6.1.1 Overall poverty

The impact of each simulated policy on overall poverty is summarized in Figure 1.

Figure 1: Change in poverty by scenario relative to baseline

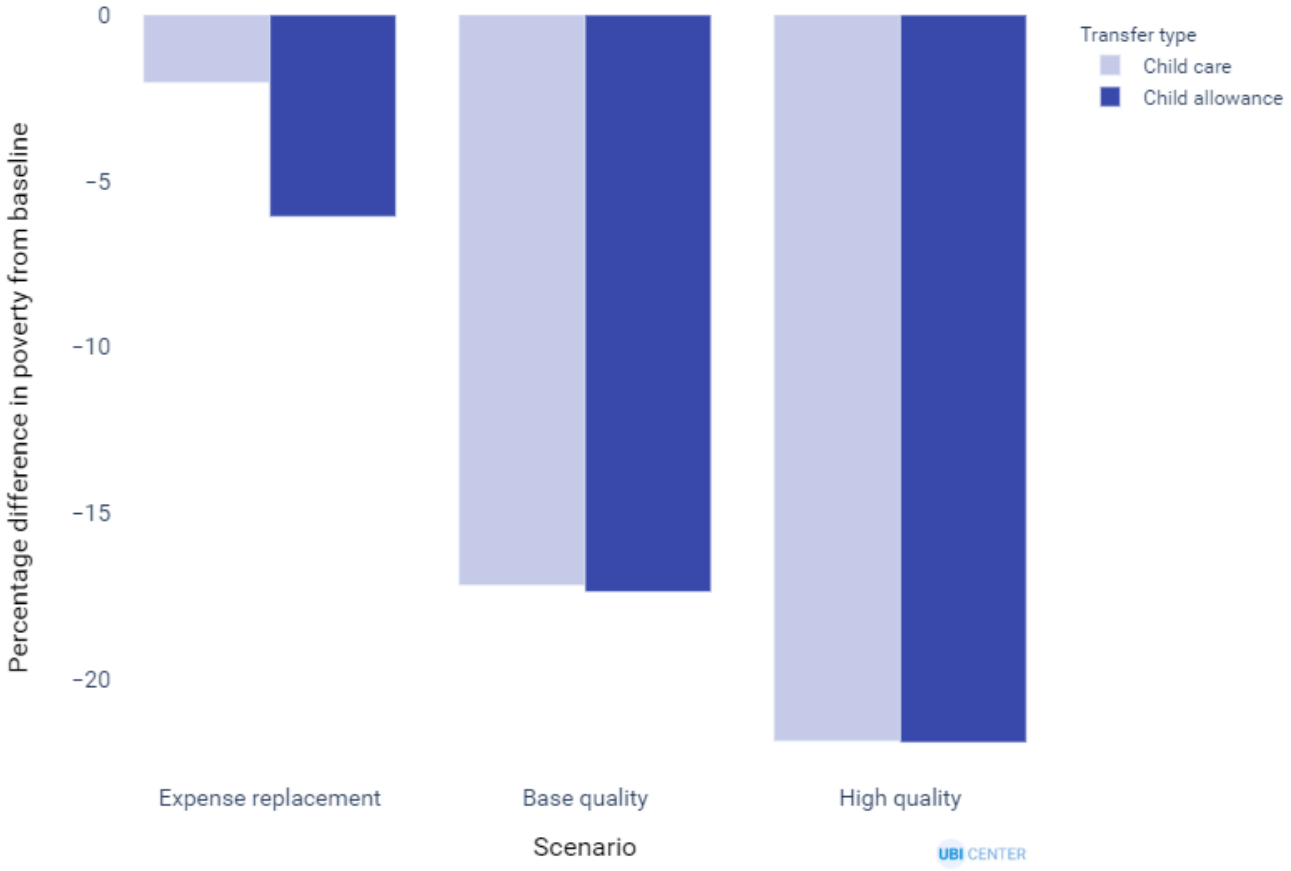

Under each of the simulated policies, poverty falls relative to baseline as the transfer to households is always positive and substantial. Under the expense replacement scenario, provision of childcare reduces poverty by 2.0%, while a child allowance reduces poverty by 6.1% from a baseline of 12.8% (0.3 and 0.8 percentage points respectively). The provision of a child allowance is therefore almost three times more effective as a tool for reducing poverty compared to a transfer equal to household expenditures on childcare.

Under the base quality scenario, both childcare and a child allowance reduce poverty by just over 17% (2.2 percentage points). This is substantially higher, commensurate with the increased size of the transfer. This reflects both the different assumptions placed on take up as well as the difference between actual expenditure on childcare compared to what is estimated to be the minimum cost of supplying childcare of a quality that meets minimum state licensing standards. We explore reasons for differences in household take up in Section 5.3.

The difference in the relative efficacy of each policy shrinks to near zero under both the low and high quality scenarios, reflecting the fact that once we assume full take up, the only differentiation between the childcare and child allowance policies is estimated state-based differences in childcare costs. These appear to have negligible implications for the efficacy of each policy in alleviating poverty. Under the high quality scenario, both policies reduce poverty

substantially further, again commensurate with the additional cost of the program. Both reduce poverty by just under 22% (2.8 percentage points).

6.1.2 Deep poverty

We turn now to deep poverty, a metric that indicates the proportion of households living under half the poverty threshold. Relative to the comparative impact of the childcare and child allowance policies in each simulation on poverty, the difference in effects is larger for deep poverty.

Figure 2: Change in deep poverty by scenario relative to baseline

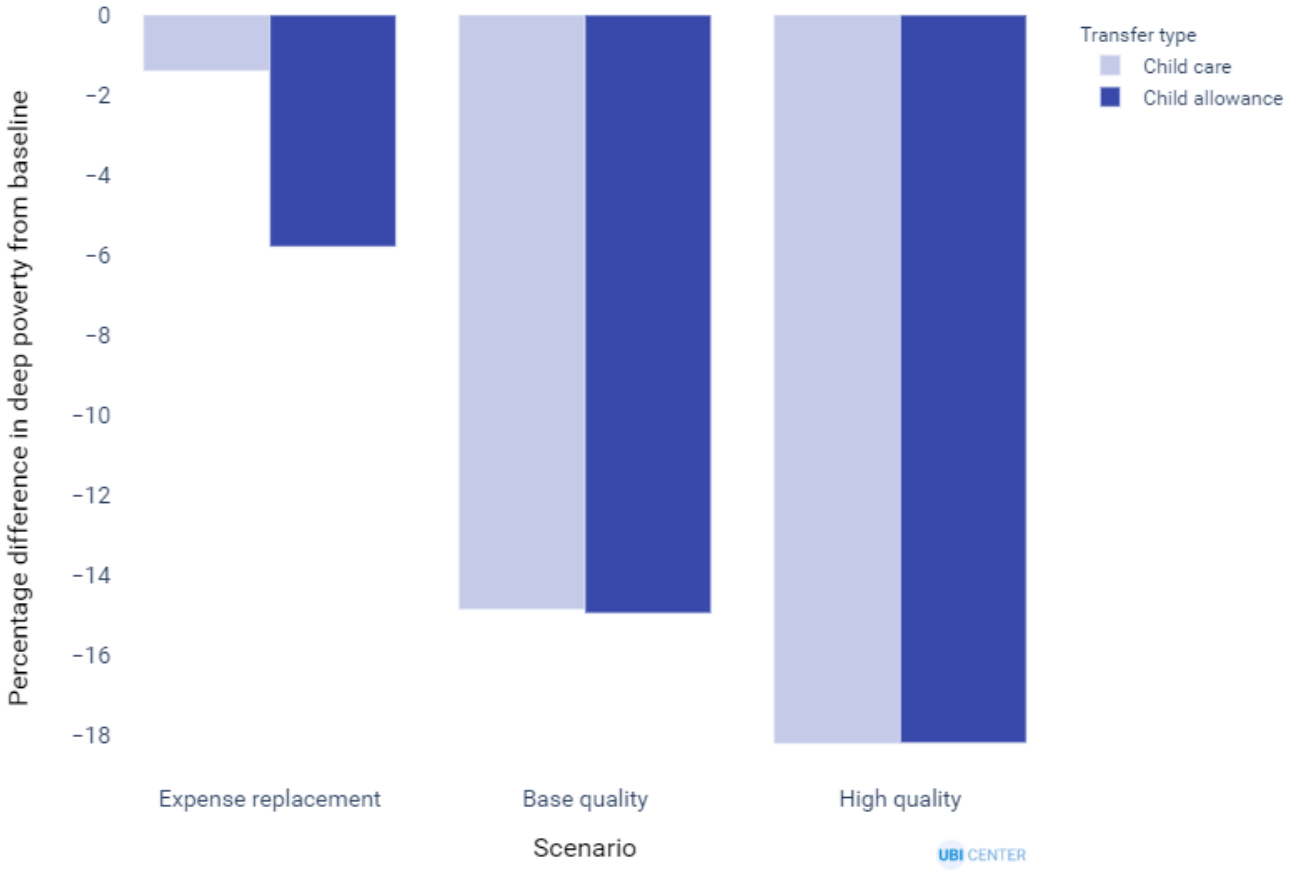

Under the expense replacement scenario, the provision of childcare causes the deep poverty rate to fall by only 1.4% (0.06 percentage points) from a baseline of 4.3%. By contrast, the child allowance causes it to fall by 5.8% (0.25 percentage points). Here the child allowance is more than four times more effective as a tool to reduce deep poverty.

Under the base and high quality scenarios, the impacts of the childcare and child allowance programs are much larger. Under the base quality scenario, both child provision and the child allowance reduce deep poverty by just under 15% (0.65 pp). Under the high quality scenario, both programs reduce deep poverty by just over 18% (-0.7 pp).

6.1.3 Child poverty

At baseline, the poverty rate for children under six years of age (12.6%) exceeds the adult poverty rate (15.0%). It is considered a priority area for policy intervention as child poverty is known to have long-term impacts numerous metrics of welfare. The impact of each program on child poverty is summarized in Table 3.

Table 5: Child poverty impacts

Scenario	**Transfer type**	**Child poverty rate**	**Percentage change**
Baseline	NA	15.0	NA
Expense replacement	childcare	13.9	-7.3
Expense replacement	Child allowance	11.5	-23.3
Base quality	childcare	4.0	-73.3
Base quality	Child allowance	3.8	-74.7
High quality	childcare	1.3	-91.3
High quality	Child allowance	1.3	-91.3

The impact of each policy on child poverty is large and there are substantial differences in impact between programs. The provision of childcare in the expense replacement scenario only reduces poverty by 7.3%, whereas a child allowance would reduce it by 23.3%. This is an important finding. If reducing child poverty is a key aim of the child allowance or provision of childcare intervention, the child allowance outperforms the childcare provision by a factor of three, even though the impacts of both programs are high.

The larger transfers associated with the base quality and high quality scenarios would reduce poverty by around 74% and 91% respectively, whether through provision for childcare or a child allowance. This is a highly consequential reduction and would make either policy a strong contender if reducing child poverty is a primary goal.

6.1.4 State-based heterogeneity

In this section, we explore the magnitude of state-based heterogeneity in the relative effectiveness of childcare provision and a child allowance in reducing poverty. We define the relative effectiveness metric as the percentage difference in transfer type efficacy by scenario, where transfer type efficacy is defined as the percentage reduction in poverty in a given state relative to baseline under a given scenario.

The effects differ markedly between scenarios. The largest differences are in the expense replacement scenario, where the child allowance outperforms childcare provision in almost every state by a substantial margin (and is never outperformed). These effects are shown in Figure 3.

Figure 3: Relative poverty impact of a child allowance compared to provision of childcare in the expense replacement scenario

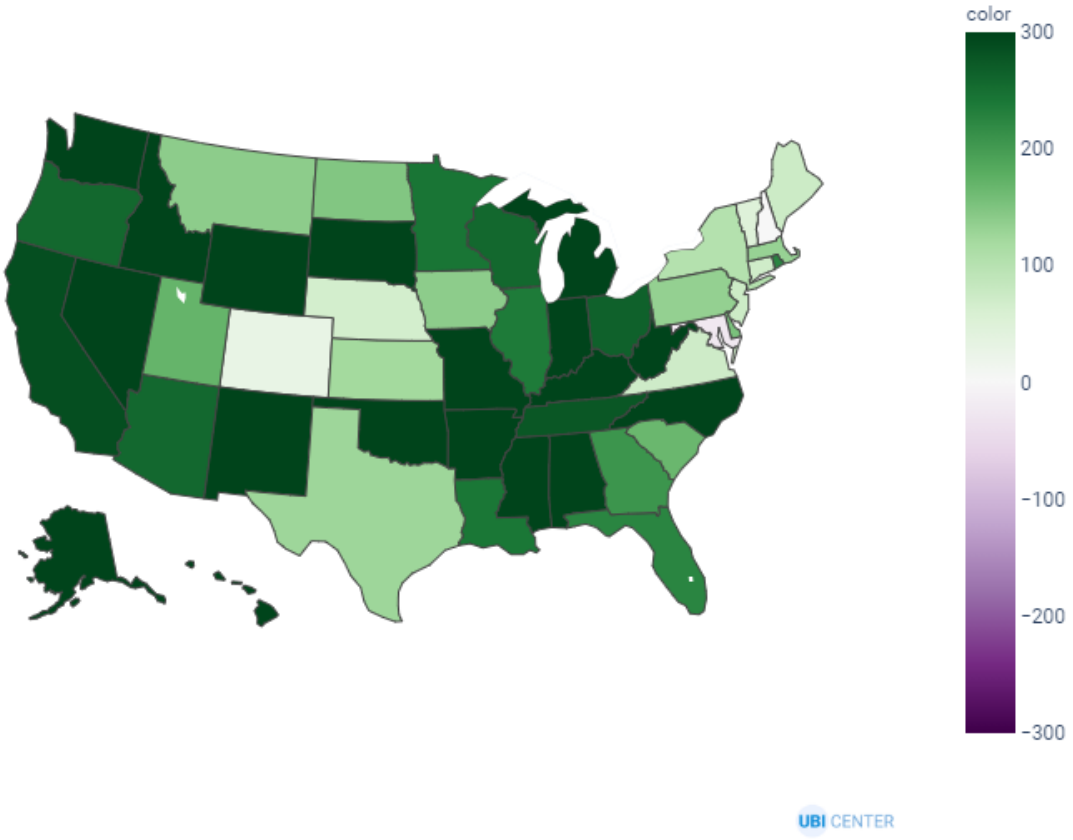

By contrast, in the base and high quality scenarios, the effects are mixed, reflecting that the overall difference in poverty impact of each transfer type under these simulations is much smaller.[5] The effects are shown in Figures 4 and 5.

[5]Note that the scale used in these figures differs from that in Figure 1. It is an order of magnitude smaller.

Figure 4: Relative poverty impact of a child allowance compared to provision of childcare in the base quality scenario

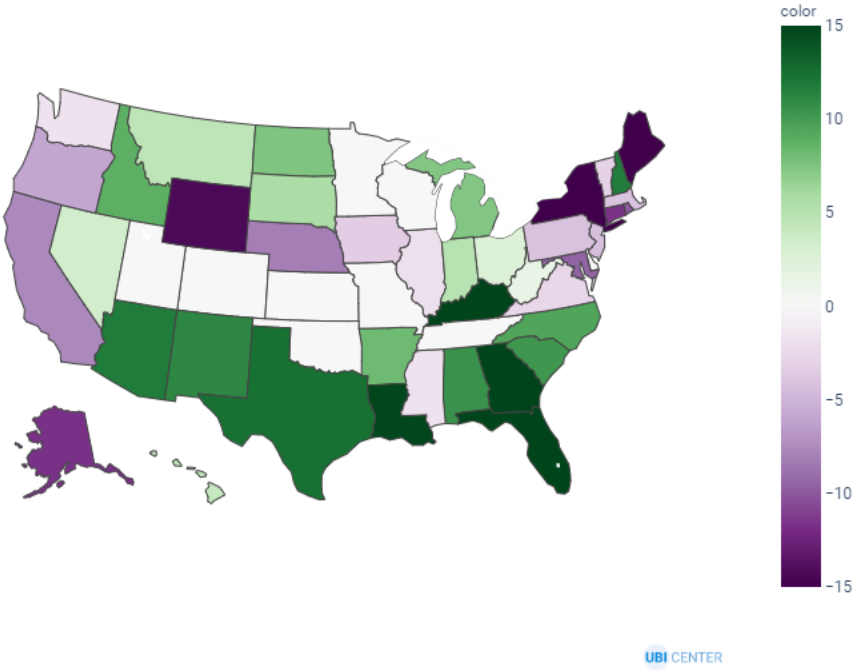

Figure 5: Relative poverty impact of a child allowance compared to provision of childcare in the high quality scenario

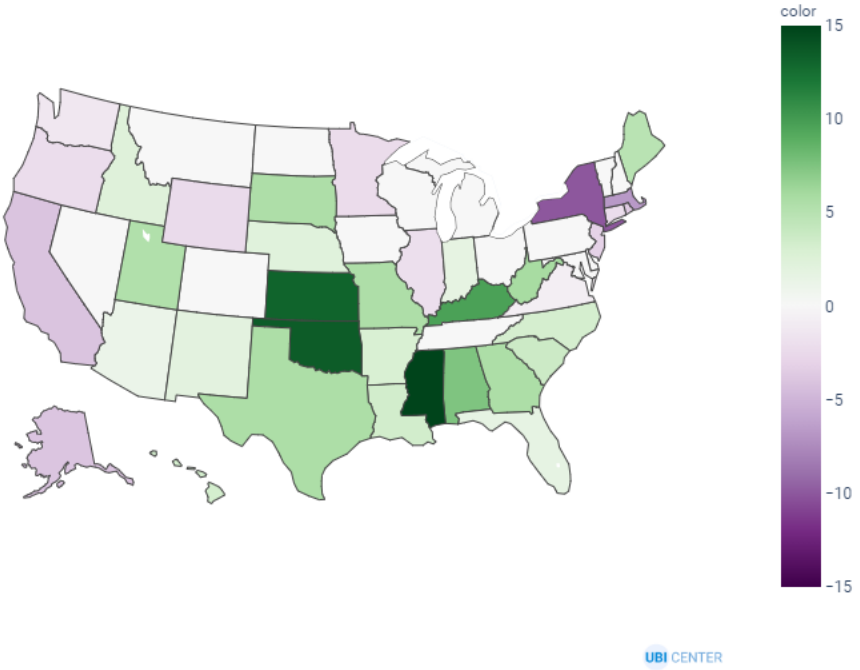

6.2 Inequality impacts

We turn now to the effect of each program on inequality. In general, the child allowance reduces the Gini coefficient by a larger amount relative to provision of childcare. Figure 3 summarizes the impact of each transfer type on the Gini coefficient by scenario.

Figure 6: Change in Gini coefficient by scenario relative to baseline

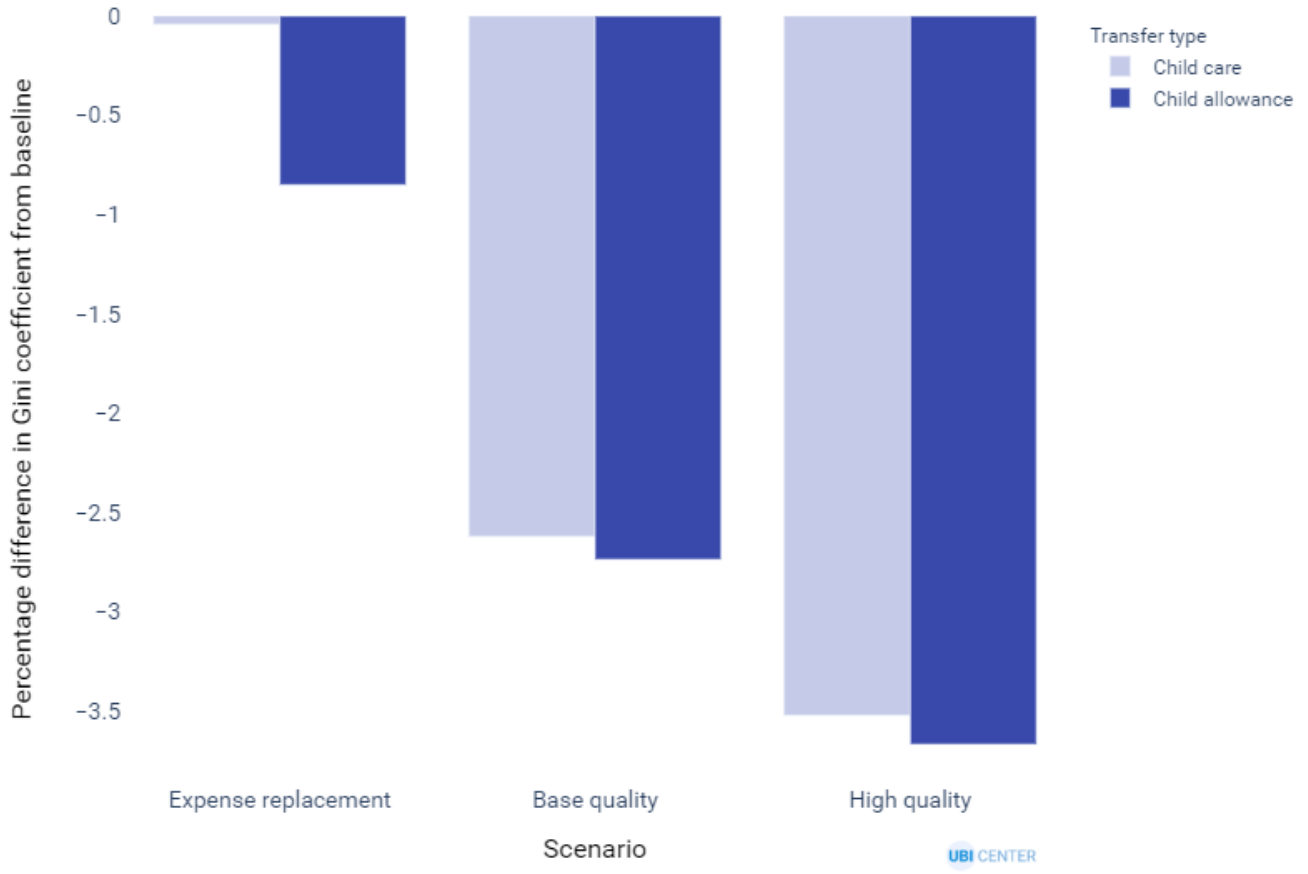

In the expense replacement scenario, provision of childcare only reduces the Gini coefficient by 0.04%, whereas the child allowance reduces it by 0.85% or roughly twenty times more.

Again, in the base and high quality scenarios, the size of the difference in impact between the childcare and child allowance shrinks substantially. In the base quality scenario, the provision of childcare reduces the Gini coefficient by 2.6%, compared to a 2.7% reduction for the child allowance. Similarly, in the high quality scenario, childcare reduces the Gini coefficient by 3.5%, whereas the child allowance reduces it by 3.7%.

7 Discussion

The results across the modelling scenarios are quite consistent - child allowances are more progressive than programs which provide childcare in their effects on poverty and inequality. This is particularly true where we assume the existing propensity to consume childcare is unchanged by the policy (in the expense replacement simulation). If poverty reduction, and particularly child poverty reduction, are key policy goals, our results provide evidence that a child allowance is a preferable policy.

In the following sections, we briefly outline some of the limitations of our simulations and opportunities for future research.

7.1 Dynamic effects

The primary limitation of the evidence we have presented is that it does not account for dynamic effects. What if guardians are sensitive to the price of childcare in their demand for childcare or their supply of labor?

In our analyses, we have assumed either that there is no new take up in response to provision of childcare or a child allowance or full take up. These two extremes bound the plausible range of values for the elasticity of demand with respect to the price of childcare offering some insight into plausible behaviours of program recipients. Of course, both scenarios are unrealistic, though it is a nontrivial task to attribute a single value to the elasticity in a dynamic model as mode of childcare and employment status of the mother are known to produce highly heterogeneous estimates[1]. An important area for future research would be a dynamic simulation that incorporates plausible elasticity estimates.

We also don't consider the labor supply decision making of parents. The average elasticity of parental employment with respect to the cost of childcare among OECD countries is around -0.34 and the elasticity of hours work with respect to childcare price is -0.34 [2]. An extension to our model would account for the dynamic decision making of guardians using elasticity estimates such as these.

7.2 Other limitations and opportunities for future research

While our study offers bounds on the impacts of various potential policies, it is likely that a policy implemented in the future will differ markedly from the scenarios that we outline here. Modelling the distributional impacts of proposed policies when details are made available should be a key concern of policy researchers looking to inform debate on the efficacy of different policies such as childcare subsidies.

Our analysis of geography-based heterogeneity in poverty impacts is also somewhat limited by the granularity of data. There is likely significant heterogeneity in childcare costs and expenditures within states. childcare expenditure is likely correlated with income due to endogenous decision making by households to live in relatively affluent areas with higher housing costs. This likely leads us to underestimate the relative impact of the child allowance compared to the provision of childcare. In future work, the use of more granular data would allow for modelling of these effects.

In this study, we have limited discussion to the immediate impacts of a child allowance or provision of childcare on measures of poverty and inequality. However, one of the key arguments of proponents of such programs is that investments in reduction of child poverty have large long-term impacts. The average child who lives in poverty is known to face worse outcomes across a broad range of outcomes relative to a child growing up in a wealthier household, and these effects persist into adulthood [3]. For example, individuals who experience childhood poverty have lower earnings, lower job market attachment and tend to have jobs in low-paying industries [4]. Studies of childcare subsidies, on the other hand, are more mixed, with some studies finding negative impacts on non-cognitive development, leading to long-term reductions in health and life satisfaction [5]. There are opportunities to simulate these effects and to address the longer term impact of the programs we consider.

8 Conclusion

We find strong evidence that provision of childcare or of a child allowance would substantially reduce poverty, reduce inequality by a small amount and reduce child poverty by a very large amount. Taken as a whole, our results provide evidence that relative to provision of childcare (or at least a transfer of equivalent size), a child allowance would be more progressive, reducing poverty and inequality by larger amounts across our simulations. That said, once we assume full take up, and only consider state-based differentiation in costs, the difference in impact between a child allowance and provision of childcare is much smaller.

Our exploratory analysis indicates that household demographics are important indicators of take up. Both take up and dynamic effects are likely important considerations for the efficacy of a child allowance or child provision program and should be the subject of further research.

References

[1] Duncan Chaplin, Philip Robins, Sandra Hofferth, Doug Wissoker, and Paul Fronstin. The price elasticity of childcare demand: A sensitivity analysis, 03 2016.

[2] Xiaodong Gong, Robert Breunig, and Anthony King. How responsive is female labour supply to child care costs - new Australian estimates. Treasury Working Papers 2010-03, The Treasury, Australian Government, April 2010.

[3] Engineering National Academies of Sciences and Medicine. *A Roadmap to Reducing Child Poverty*. The National Academies Press, Washington, DC, 2019.

[4] Rune V. Lesner. The long-term effect of childhood poverty. *Journal of Population Economics*, 31(3):969–1004, July 2018.

[5] Michael Baker, Jonathan Gruber, and Kevin Milligan. The long-run impacts of a universal child care program. *American Economic Journal: Economic Policy*, 11(3):1–26, August 2019.