Optimal Transfers with Children's Utility

Lea Bart¹ and Emily $Horton^2$

¹University of Michigan

²U.S. Department of the Treasury

November 2024

This paper studies the optimal tax and transfer system when the social planner's objective is to maximize utility of all individuals, including children, and they can use household composition as a fixed tag. While existing research has estimated optimal transfer schedules, it has typically done so only in the context of adults, or without considering children's utility as part of the planner's objective function. We simulate the optimal tax and transfer system for unmarried households accounting for children's consumption, adults' consumption, and adults' labor supply responses on the intensive and extensive margin. We place particular emphasis on estimating optimal transfers to non-earning households of different compositions. Our findings demonstrate that the level of children's consumption needs relative to adults is of first order importance for determining optimal transfers to single parent households. We find that within most income groups, current policy places minimal weight on redistribution from childless households to households with children. This is inconsistent with standard government assumptions regarding families' resource needs.

We thank Ashley Craig, Jim Hines, Joel Slemrod, Sarah Miller, and the participants in the University of Michigan Public Finance Workshop for helpful comments and suggestions. Bart's work was conducted with support from an NICHD training grant to the Population Studies Center at the University of Michigan (T32 HD007339). This research was conducted while Horton was a graduate student at the University of Michigan and an employee at the U.S. Department of the Treasury. The findings, interpretations, errors, and conclusions expressed in this paper are entirely those of the authors and do net necessarily reflect the views or official positions of the U.S. Department of the Treasury or the Office of Tax Analysis. Any taxpayer data used in this research was kept in a secured Treasury or IRS data repository, and all results have been reviewed to ensure that no confidential information is disclosed.

1 Introduction

United States transfer policy places a substantial emphasis on redistribution to children. Many transfer programs are only available to families with children (such as Temporary Assistance for Needy Families, the Supplemental Nutrition Program for Women, Infants, and Children, and the Child Tax Credit) or are more generous to children (such as Medicaid and the Earned Income Tax Credit). However, transfers to the poorest children are moderated by concerns about discouraging the labor force participation of their parents. These concerns have led to many programs targeted to children and families being contingent on employment or earnings, leaving children with non-working parents with limited transfers (Aizer, Hoynes, and Lleras-Muney 2022).¹

We develop an optimal tax model that balances children's consumption needs against distortions in parents' labor force participation in the context of single adult households. While optimal tax models are designed to quantify the tradeoff between redistribution and labor supply distortions, they typically focus on redistribution between adults and abstract away from redistribution to children. We calibrate the consumption needs of single parent households relative to single childless adults using equivalence scales, which are widely used in policy to compare well-being across different family types. We include both extensive margin and intensive margin labor supply responses in the model, which allows us to estimate optimal transfers to households without earnings.

Including children's consumption needs substantially increases optimal transfers to families with children. Under our preferred specification, which uses conservative assumptions, we estimate that the optimal transfer to single parents without earnings is roughly \$29,500 when children's resource needs are included in the planner's problem, compared to just \$17,100 when those needs are not considered. Assumptions about the value of children's

¹For example, Temporary Assistance for Needy Families (TANF) usually requires work to receive benefits, and the Earned Income Tax Credit (EITC) and Child Tax Credit (CTC) are unavailable to families without earnings. For a detailed analysis of shifts in the demographic distribution of benefits over time, see Moffitt (2015).

resource needs relative to adults are of first order importance. For example, moving from the U.S. Census Bureau's assumption that children consume roughly half as much as adults to the Congressional Budget Office's assumption that children consume an equal amount to adults increases the optimal transfer to single parent households without earnings by \$8,700 (29.5%). In comparison, reducing single parents' extensive margin labor supply elasticity from the high to low end of estimates from the literature (0.5 to 0.3) only increases the optimal transfer to single parents without earnings by \$2,700 (4.4%).

We find that even if the policymaker does not highly value redistribution between income groups, it is generally still optimal to transfer substantial resources to low-income single parents. This is because there are both equity and efficiency justifications for transferring resources to these households. The equity motive is driven by the fact that low-income single parent households have minimal resources and multiple people consuming them. The efficiency motive is driven by the fact that families have economies of scale in consumption; because they share some goods, such as housing, they generate more than a dollar of consumption per dollar of income.

We also use the model to infer how much the current tax system values redistribution from higher-income households to lower-income households and from childless households to single parent households. To do so, we calculate the inverse-optimum welfare weights that rationalize the current tax and transfer system (Bourguignon and Spadaro 2012). We find that the current system places a relatively low value on redistribution from higher-income to lower-income households.² Within an income group, the current tax and transfer system generally places very little weight on redistribution from childless households to single parent households. The exception is low-income working households; within this income group, current policy values single parent households' marginal consumption almost twice as much as their childless counterparts.

This paper contributes to an extensive literature on the importance of children in optimal

 $^{^2{\}rm This}$ is consistent with the findings of Bargain et al. (2014), Lockwood and Weinzierl (2016), and Hendren (2020).

taxation. Prior work has motivated redistribution to families with children by considering the impact of parental income and investments on their future ability (Cigno and Pettini 2002; Cigno, Luporini, and Pettini 2003, 2004; Casarico, Micheletto, and Sommacal 2015; Gelber and Weinzierl 2016), modeling children as consumption goods in parents' utility (Kurnaz 2021), using family size as a tag for parental ability (Balestrino, Cigno, and Pettini 2002; Cigno, Luporini, and Pettini 2004; Blumkin, Margalioth, and Sadka 2015), and allowing childcare costs to impact parental consumption and labor supply (Blundell and Shephard 2012; Apps and Rees 2018; Guner, Kaygusuz, and Ventura 2020; Bastani, Blomquist, and Micheletto 2020).

In contrast, redistribution in our model occurs because the policymaker directly values adults' *and* children's utility from consumption. Under this assumption, the policymaker's primary motive for redistribution to children is providing resources to individuals who cannot provide for themselves. Conceptually, this paper relates most closely to Cremer, Dellis, and Pestieau (2003), which allows altruistic parents to allocate consumption to children under a simplified set of assumptions about family utility, labor supply responses, and the structure of the tax system.

We contribute to the optimal tax literature by extrapolating children's consumption, utility, and marginal utility from equivalence scales. To do this, we build on the assumptions used in the economics literature to estimate equivalence scales from data on household expenditures (Browning, Chiappori, and Lewbel 2013; Dunbar, Lewbel, and Pendakur 2013; Calvi et al. 2023). We implement the resulting family utility function in an optimal tax model with extensive margin labor supply responses, generalized from Saez (2002).³ The advantage of deriving family utility from equivalence scales is that it allows us to model the optimal tax and transfer system under a set of transparent assumptions that are consistent with the assumptions used by government agencies in policy-making and poverty measurement.

 $^{^{3}}$ The literature on optimal taxation with extensive margin labor supply responses has been further developed by Laroque (2005), Choné and Laroque (2010), and Jacquet, Lehmann, and Van der Linden (2013), among others.

Overall, our results highlight the importance of considering children's resource needs when designing optimal transfer policy. Ignoring these needs leads to significantly lower transfers to families than is optimal, even when overall preferences for redistribution are low. Current tax and transfer policy places minimal weight on children in most income groups, which is inconsistent with government estimates of children's resource needs.

2 Model

2.1 Model Framework

We develop an optimal tax model that distinguishes between households of different sizes and allows adults to adjust their labor supply along both the extensive and intensive margins in response to the tax schedule. Family size is treated as a fixed tag, such that the government can observe the number of children in the household and set different tax/transfer schedules for families of different sizes.⁴ We model households as being run by a single, altruistic decision maker – for this reason, we focus on unmarried adults. Adults vary in skill (represented by wages), fixed costs of work, and the number of children in their household. Following the Saez (2002) framework, adults choose from a discrete set of possible earnings levels, each of which is associated with a net tax or transfer amount set by the social planner.

2.2 Household Problem

The relevant unit in this paper is the family, consisting of individual members who are either adults (A), or children (K). The utility of the family is the sum of the utilities of

⁴In general, there is limited empirical evidence that low-income women in the United States adjust their childbearing in response to transfer programs, even those explicitly aimed at reducing births (Horvath-Rose, Peters, and Sabia 2008; Kearney 2004; Joyce et al. 2004; Baughman and Dickert-Conlin 2009). Furthermore, qualitative research suggests that childbearing among low-income women is primarily driven by non-economic factors (Edin and Kefalas 2005). There is some evidence that multifamily households adjust which taxpayer claims dependent children in response to tax incentives (Jones and O'Hara 2017; Splinter, Larrimore, and Mortenson 2017). Accounting for this margin of behavioral response would slightly increase the cost of increasing transfers to families with children.

consumption of each of its members, minus a disutility of earning income:

$$u_F = u_A(c_A) + K u_K(c_K) - v_{w,x,K}(z)$$
(1)

K is the total number of children in the household. u_A is the utility of the single adult as a function of their consumption c_A . We assume each child in the household has utility u_K , which is a function of their consumption c_K . The utility cost for the family to earn income z is represented by $v_{w,x,K}(z)$ and will vary with the adult's skill level w, fixed cost of work x, and number of children K.

We assume that family size is exogenous, so the household takes K as given. In this set up, children are included in the social welfare function through the utility the children themselves receive from consumption. This underpins the theoretical justification for redistribution toward larger families on the basis of larger families having higher consumption needs. This is in contrast to models with endogenous fertility, where transfers based on children are justified with an externality or equity argument about the utility of current adults (see Cigno, Pestieau, and Rees (2011) for an overview).

We assume a unitary model of the household, where the household is run by a single, altruistic decision maker who maximizes the total utility of all members. This generates the following first order condition for the optimal allocation of consumption within the household:

$$u'_{A}(c^{*}_{A}) = u'_{K}(c^{*}_{K}) \tag{2}$$

We limit our analysis to single adult households, where the assumptions of altruistic decisionmakers and optimal allocation of consumption are more plausible. Empirical evidence suggests that the allocation of consumption within married households generally varies with the bargaining power of the spouses, which contradicts the unitary model of the household (see Chiappori and Meghir (2015) for an overview). Without these assumptions, single parent households would be less efficient at generating family utility from income, which would decrease the social value of transfers to families over childless adults relative to our model.

We allow households to take advantage of economies of scale in consumption. The household chooses c_A , c_K , and z to maximize family utility subject to the budget constraint

$$(1 + \alpha_K)I = c_A + Kc_K \tag{3}$$

where
$$I = z - T_{K,z}$$

The household's after-tax income I is defined as its earnings z less its tax liability $T_{K,z}$. The scaling factor α_K represents the household's economies of scale in consumption. Because some consumption goods are shared between family members, the sum of individual consumption done by each family member, $c_A + Kc_K$, can exceed the family's after-tax income I by a factor of α_K . We assume that economies of scale α_K are a function of family size but not income, following the equivalence scale literature (Lewbel and Pendakur 2008; Calvi et al. 2023).⁵

Finally, we assume that the optimal allocation of consumption within the household does not vary by income, so each child's consumption allocation can be represented as a fixed proportion p of the adult's consumption allocation:

$$c_K^* = p c_A^* \tag{4}$$

The assumption that the share of resources allocated to each household member are independent of total household resources is another standard assumption in the equivalence scale literature (Dunbar, Lewbel, and Pendakur 2013; Calvi et al. 2023). This assumption is supported by empirical evidence in Betson (1996) and Menon, Pendakur, and Perali (2012),

⁵Put differently, we assume the household takes its economies of scale as exogenous. While this is likely sufficient to capture broad differences in economies of scale across household sizes, note that Chiappori (2016) argues this assumption is an oversimplification for the purpose of identifying equivalence scales from demand data. This is because economies of scale reflect an endogenous choice of how to allocate spending across a variety of shared and private goods.

which find that the share of income spent on children does not vary significantly between families of the same size with different income levels. The per-child consumption allocation p does not vary by family size, because there is no reason why each individual child's resource needs relative to an adult's should depend on family size.⁶

This assumption that children are allocated a fixed proportion of adult consumption constrains the functional form of the child's utility function. It also simplifies the household problem into a two-step decision, where the parent first chooses how to allocate consumption between members of the household and then chooses how much income to earn. Applying equations 2 and 4 and integrating allows us to rewrite the child's utility of consumption u_K as a scaled version of the adult's utility of consumption u_A (plus or minus an additive constant, not shown):

$$u_K(c_K) = p u_A\left(\frac{c_K}{p}\right) \tag{5}$$

Using the consumption allocation from equation 4 and the functional form of children's utility from equation 5, we can write the household's utility as a function of its earnings z and taxes T.

$$u_F = (1 + pK)u_A \left(\frac{1 + \alpha_K}{1 + pK}I\right) - v_{w,x,K}(z)$$
(6)
where $I = z - T_{K,z}$

As will be discussed in the next section, the key input to the social planner's problem is the household's marginal utility of income. This measures the marginal utility of increasing the household's after-tax income I by decreasing its taxes T. Differentiating equation 6 yields the household's marginal utility of income:

⁶In a collective model of the household, children in larger families could have different bargaining power than children in smaller families. In this case, children in larger families may be allocated different resources than children in smaller families, even though their underlying resource needs are the same.

$$MUI = (1 + \alpha_K)u'_A\left(\frac{1 + \alpha_K}{1 + pK}I\right)$$
(7)

In the case of a childless adult $(K = 0 \text{ and } \alpha_K = 0)$, the marginal utility of income is simply the marginal utility of consumption:

$$MUI_{K=0} = u'_A(I) \tag{8}$$

For both childless adults and single parent households, the marginal utility of income is a function of the household's after-tax income. Because the utility cost of earning income is additively separable from the utility of consumption, the marginal utility of income is not a function of wages or fixed costs of work.

2.3 Social Planner Problem

We extend the social planner's problem from Saez (2002) to allow for multiple types of families. The tax schedule T is a set of net tax levels $T_{z,k}$ for each discrete earnings level zand family composition k, where k is the number of children and family size equals k + 1. The social planner chooses a tax schedule to maximize a utilitarian social welfare function given a fixed revenue constraint:

$$\max_{T} \sum_{k} \int_{w} \int_{x} U_{w,x}(z(T), T_{z,k}, k) f_{k,w,x} dx dw$$
(9)

$$\sum_{k}\sum_{z}h_{z,k}(T)T_{z,k} = H \tag{10}$$

Households vary by number of children k, wages w, and fixed costs of work x. The distribution of underlying types is represented by f. Each household's utility is a function of their earnings z, number of children k, tax amount $T_{z,k}$, and costs of earning income (which depend on wages w and fixed costs of work x). Based on the tax schedule T, households

choose a discrete earnings level z. This generates a realized earnings distribution h, which is a function of the tax schedule. Under this realized earnings distribution, the government must raise an average revenue of H per household.

The social planner's first order condition with respect to the tax amount at a given earnings level Z and number of children K is

$$\frac{d\mathcal{L}}{dT_{Z,K}}: 0 = \frac{dU_{Z,K}}{dT_{Z,K}} h_{Z,K} + \lambda \left[\sum_{z} \frac{dh_{z,K}}{dT_{Z,K}} T_{z,K} + h_{Z,K} \right]$$
(11)

Alternatively, the first order condition can be written as

$$g_{Z,K} = 1 - \frac{1}{h_{Z,K}} \sum_{z} \frac{dh_{z,K}}{d(Z - T_{Z,K})} T_{z,K}$$
(12)
where $g_{Z,K} = \frac{MU_{Z,K}}{\lambda}$

The left hand size of equation 12 represents the utility gain from increasing the after-tax incomes of households with K kids earning income Z, scaled to a dollar value by the marginal value of public funds λ . The right hand size of the equation represents the revenue loss from decreasing $T_{Z,K}$, which is a combination of the mechanical effect of tax changes and the fiscal externality from labor supply responses. At the optimal tax schedule, the utility gain from a tax decrease on any group should be exactly offset by the revenue loss.

The fiscal externality generated by labor supply responses to changes in the tax schedule appears in the $\frac{dh}{dT}T$ terms and reflects how revenue changes as the realized income distribution h changes. To make the fiscal externality term concrete, we assume a family in earning group z under the current tax system has limited options in their choice of work. They can move into an adjacent earning group z + 1 or z - 1 (intensive margin response) or can leave the labor force (extensive margin response). Under this assumption, for a given number of kids K, the fraction of the population in earning group z is not a function of the entire tax schedule – it is only a function of $T_{z,K}$, $T_{z-1,K}$, $T_{z+1,K}$ and $T_{0,K}$. This simplifying assumption allows us to define mobility elasticities, which represent how the earnings distribution changes with respect to a change in the tax schedule:

$$\eta_{z,k} = \frac{dh_{z,k}}{d(I_{z,k} - I_{0,k})} \frac{I_{z,k} - I_{0,k}}{h_{z,k}}$$
(13)

$$\zeta_{z,k} = \frac{dh_{z,k}}{d(I_{z,k} - I_{z-1,k})} \frac{I_{z,k} - I_{z-1,k}}{h_{z,k}}$$
(14)

Empirically, these terms are derived from intensive and extensive margin labor supply elasticities.⁷ The intensive margin mobility elasticity determines the share of earnings group z that will move down to income group z - 1 when T_z is increased. The extensive margin mobility elasticity determines the share of earnings group z that will leave the labor force and move to zero earnings when T_z is increased.⁸

Applying these assumptions to the first order conditions and rearranging, we get that at the optimal tax schedule, for all combinations of Z and K,

$$\frac{T_{Z,K} - T_{Z-1,K}}{I_{Z,K} - I_{Z-1,K}} = \frac{1}{\zeta_{Z,K} h_{Z,K}} \sum_{z} h_{z,K} \left[\left(1 - \frac{MUI_{z,K}}{MVPF} \right) - \eta_{z,K} \frac{T_{z,K} - T_{0,K}}{I_{z,K} - I_{0,K}} \right] \text{ for } Z > 0 \quad (15)$$

and

$$(1 - g_{0,K})h_{0,K} = -\sum_{z} \eta_{z,K} h_{z,K} \frac{T_{z,K} - T_{0,K}}{I_{z,K} - I_{0,K}} \text{ for } Z = 0$$
(16)

These first order conditions are identical to those in Saez (2002), except everything is now indexed by the number of kids K. The earnings distribution, labor supply elasticities, and marginal utility of income are all allowed to vary by the number of children K. The main

⁷The intensive margin elasticity ζ in our model is slightly different from the classic earnings elasticity ϵ . In the model, we translate between the two such that $\zeta_z(I_z-I_{z-1}) = \epsilon_z I_z$, following Saez (2002).

⁸Parameterizing extensive margin labor supply responses in this way focuses on substitution effects (responses to changes in the return to working $I_z - I_0$) and abstracts away from income effects (responses to changes in the level of I_0 , holding constant the return to working). This implies that the introduction of an unconditional transfer payment would have no effect on labor supply. Recent evidence on the effect of unconditional transfers on labor supply is mixed (Feinberg and Kuehn 2018; Jones and Marinescu 2022; Lippold and Luczywek 2023; Baker, Messacar, and Stabile 2023; Vivalt et al. 2024; Sauval et al. 2024). To the extent that income effects are present, they would increase the distortions generated by increasing transfers.

difference between this model and Saez (2002) comes from the marginal utility of income term. In this model, the marginal utility of income will not only influence redistribution between families with different earnings levels, but also between different family types. As the marginal utility of income for a family with children increases relative to the marginal utility of income for a childless adult, the planner will prefer larger transfers to families with children.

3 Simulation of the Optimal Tax System

3.1 Model Calibration

3.1.1 Equivalence Scale

We apply the interpretation of equivalence scales used by Betson (1990) and formalized as an "indifference scale" by Browning, Chiappori, and Lewbel (2013).⁹ The equivalence scale s(K) is defined such that an adult in a household with K kids and income y would need income $\frac{y}{s(K)}$ to achieve the same standard of living as a single adult living alone.¹⁰ We interpret this to mean an adult with K kids and income y consumes

$$c_A = \frac{y}{s(K)} \tag{17}$$

For example, the U.S. Census Bureau's equivalence scale for a family with one adult and two children is s(2) = 1.79. This means that an adult with two children with an income of \$30,000 consumes $\frac{\$30,000}{1.79} = \$16,760$, which gives them the same standard of living as a single childless adult with an income of \$16,760.

⁹Browning, Chiappori, and Lewbel (2013) study a setting where household members demand individual goods, rather than a single composite consumption good. Chiappori (2016) provides an intuitive overview of indifference scales.

¹⁰We interpret "standard of living" to mean utility from consumption, abstracting away from any disutility from earning income. This is because equivalence scales don't adjust for whether income comes from labor or from non-labor income.

We calibrate the model using the functional form of equivalence scales most commonly used by policymakers. For single adult families, these equivalence scales generally take the form:

$$s(K) = \begin{cases} 1 & K = 0\\ (1+pK)^F & K \ge 1 \end{cases}$$
(18)

Where p represents the resource needs of a child relative to an adult and F represents the economies of scale in consumption for a family, where F = 1 denotes no economies of scale and F = 0 denotes perfect economies of scale. We calibrate the parameter p in the model to match the parameter p in the equivalence scale, as these both conceptually represent the family's expenditures on a child relative to an adult (Betson 1996).

We calibrate the economies of scale parameter α_K in the model by equating the adult's level of consumption from the household problem with the adult's level of consumption from the definition of the equivalence scale

$$c_A = \frac{1 + \alpha_K}{1 + pK} I = \frac{I}{s(K)} \tag{19}$$

This gives

$$1 + \alpha_K = \frac{1 + pK}{s(K)} \tag{20}$$

For our baseline model we use a simplified version of the equivalence scales used by the U.S. Census Bureau, which originates in Betson (1996). Betson estimates p and F using data on how much households spend on adults and children across household sizes and income levels. Our baseline model adopts a simplified version of the parameters used by Census: p = 0.5, and F = 0.7 (Short 2011). The actual equivalence scale used by Census assumes the first child in a single parent family has 0.8 times the consumption of an adult, while all other

children have 0.5 times the consumption of an adult.¹¹ We simplify the scale so that each additional child has identical resource needs. Using our simplified scale yields s(2) = 1.62 compared to the full version, which yields s(2) = 1.79. The simplification thus reduces the incentive to transfer to families relative to the full Census version.

3.1.2 Marginal Utility of After-Tax Income

The marginal utility of after-tax income for families depends on the functional form of the adult's utility function. We simulate the model using a CRRA utility function for adults. The household's marginal utility thus takes the form

$$MUI_F = \frac{1+pK}{(1+pK)^F} \left(\frac{I}{(1+pK)^F}\right)^{-\theta}$$
(21)

In our utilitarian set up, the redistributive preferences of the social planner are entirely driven by the concavity of the utility function, such that θ drives the amount of redistribution the government wants to do. Our baseline model assumes the adult's utility function has a concavity of $\theta = 1$, which corresponds to the conventionally used logarithmic benchmark. Inverse optimum exercises suggest that $\theta = 1$ was a fair representation of redistributive preferences prior to the Tax Reform Act of 1986, although more recently the tax system has implied lower redistributive preferences (Lockwood and Weinzierl 2016).

Note that the family's marginal utility of income is increasing in p. If children have higher resource needs, the social planner has more incentive to redistribute to them. However, the family's marginal utility of income is not necessarily increasing in its economies of scale F. Intuitively, economies of scale affect the family's marginal utility in two offsetting ways they allow the family to consume more at a given income (which decreases their marginal utility of income) and they increase the consumption generated by giving the family an extra dollar of income (which increases their marginal utility of income). When $\theta < 1$, the former

¹¹This appears to be an attempt to reflect differences in economies of scale between single and married families, rather than an indication that there are actual differences in consumption needs between the first child and subsequent children.

dominates, and the family's marginal utility of income is decreasing in F. When $\theta > 1$, the latter dominates, and the family's marginal utility of income is increasing in F. Under our baseline case of log utility ($\theta = 1$), these exactly offset each other, and the family's marginal utility of income does not depend on F.

It is useful to consider how a family's marginal utility of income compares to a single adult with the same income. The ratio of a family's marginal utility to a single adult's marginal utility at a given income level I can be expressed as

$$\frac{MUI_F}{MUI_S} = (1 + pK)^{1 - (1 - \theta)F}$$
(22)

Under the baseline specification (p = 0.5, F = 0.7, $\theta = 1$, K = 2), $\frac{MUI_F}{MUI_S} = 2$. This means that a single parent household with two children would have double the marginal utility of a single adult earning the same income.

Note that $\frac{MUI_F}{MUI_S}$ is a function of the CRRA utility parameter θ in addition to the equivalence scale parameters p and F. The parameter θ represents the welfare cost of consumption inequality between individuals. In standard models, θ governs how much the social planner values redistribution from high-income to low-income households. In this model, θ also influences how much the social planner values redistribution from smaller households to larger households within a given income level. This is because smaller households have higher income and consumption per person than larger households at the same income level.

3.1.3 Labor supply responses

We choose labor supply elasticity estimates near the upper bound of the empirical literature. Because we allow transfers to generate relatively large labor supply distortions, this assumption will generate conservative estimates of optimal transfers. We assume that only those earning \$30,000 or less respond along the extensive margin, and that those with higher earnings have extensive margin elasticities of 0. (See Appendix Figure A4 for robustness to this assumption.)

We use the literature regarding single mother responses to expansion of the Earned Income Tax Credit to calibrate the extensive margin labor supply elasticity for single parents (Hoynes and Patel 2018; Bastian and Jones 2021; Bastian 2020; Eissa and Liebman 1996; Meyer and Rosenbaum 2001; Chetty, Friedman, and Saez 2013). The upper end of the range of elasticity estimates in the literature is roughly 0.5, which is what we use in our baseline model.¹² We calibrate the extensive margin elasticity of childless adults to be 0.3 based on the results of the Paycheck Protection RCT, which offered childless adults an earnings bonus equivalent to an EITC expansion (Miller et al. 2018).

For the intensive margin elasticity we use the unified estimate from Chetty (2012), 0.33, and assume it is identical for single parents and childless adults. Many studies of individual tax reforms, including those focused on the EITC, estimate intensive margin responses well below this level. It is possible that responses to smaller reforms are dampened by optimization frictions and that larger reforms would elicit larger responses, especially over the long term. We prefer Chetty's estimate for our baseline specification because it scales up intensive margin responses to account for optimization frictions.

3.1.4 Initial conditions

Initializing the model requires three more objects: the current earnings distribution, the tax schedule that generated the current earnings distribution, and the target average revenue per single adult household.

Table 1 shows the current earnings distribution used to calibrate the model. We measure

¹²Hoynes and Patel (2018) study how single mothers responded to the EITC expansions in the 1990s and estimate an extensive margin elasticity of 0.14-0.35 (see online addendum). Bastian and Jones (2021) pool variation from a broad range of EITC expansions over time and estimate an extensive margin elasticity of 0.33. Chetty, Friedman, and Saez (2013) calculate extensive margin elasticities implied by earlier studies. They find that the Eissa and Liebman (1996) study of the EITC expansion in 1986 implies an extensive margin elasticity of 0.30, and the Meyer and Rosenbaum (2001) study of EITC expansions and welfare cuts in the 1990s implies an extensive margin elasticity of 0.43. Bastian (2020) estimates single mothers' response to the introduction of the EITC in the 1970s and estimates an extensive margin elasticity of 0.58. Because labor supply elasticities were likely larger in the 1970s, an elasticity of 0.5 is still a conservative estimate of current labor supply responses.

the earnings distribution based on individual earnings (wage and salary, self-employment, and farm income) of unmarried adults in the pooled 2018-2020 March Current Population Survey (CPS), excluding students and those over age $60.^{13}$ We discretize the earnings distribution by dividing households into 14 earnings groups and two size groups – those without children and those with children. Children are identified from the number of the adult's own children living in the same household as them. Single parent families have an average of 1.8 children each, so we calibrate the marginal utilities of single parent families based on having two children (K = 2).

We smooth the earnings distribution observed in the CPS because the CPS earnings distribution is heaped at round numbers (Schwabish 2007). Within each household type (single adults without children and single parents with children), we estimate a kernel smoothed probability density function on all households with positive earnings.¹⁴ We supplement the smoothed earnings distribution with the discrete mass of workers with zero earnings reported in the CPS. Appendix Figure A1 shows that this smoothed earnings distribution closely approximates the CPS earnings distribution.

We represent the current tax and transfer system for each household type using a linear approximation of the current tax system. Section 6.1 describes the current tax system in detail using CPS data on taxes paid and transfer benefits received by single adult households. A linear fit of the current tax system yields a lump sum transfer of \$7,685 with a marginal rate of 32% for childless adults and a lump sum transfer of \$13,685 with a marginal rate of 28% for single parents. Appendix Figure A2 shows that the simulation results are robust to calibrating the model using the more detailed approximation of the current tax system from Section 6.1.

The target government revenue (H) is calibrated to \$3,430 per household under the

¹³We exclude adults that report negative earnings or that report investment income (interest, dividends, and rents) that is more than a quarter of their earned income. We exclude these groups because their earnings are not a good proxy of their available income.

¹⁴We choose an Epanechnikov kernel with a bandwidth of \$10,000 because CPS earnings are heaped at intervals of \$5,000.

Lower earnings	Representati	ve earnings	Share of households		
threshold	Childless adults	Single parents	Childless adults	Single parents	
0	0	0	0.159	0.052	
5,000	$7,\!662$	$7,\!673$	0.032	0.011	
10,000	$12,\!626$	$12,\!634$	0.040	0.014	
15,000	$17,\!593$	$17,\!599$	0.046	0.016	
20,000	22,964	22,967	0.057	0.020	
$25,\!831$	$29,\!579$	29,567	0.074	0.026	
33,362	38,053	38,014	0.086	0.030	
43,089	48,950	48,880	0.085	0.027	
$55,\!651$	$63,\!009$	62,938	0.069	0.021	
$71,\!876$	81,168	$81,\!087$	0.047	0.014	
92,832	104,320	$104,\!278$	0.028	0.008	
$119,\!897$	$135,\!107$	135,223	0.014	0.004	
154,853	$173,\!555$	173,732	0.007	0.002	
200,000	403,713	$393,\!013$	0.008	0.002	
Total			0.753	0.248	

Table 1: Initial income distribution

Notes: Earnings distribution estimates are from pooling the 2018-2020 March CPS. The sample is unmarried adults, restricted to those age 18-60 and excluding students. Workers with negative earnings or with investment income greater than a quarter of earned income are excluded from the sample. The left column shows cutoff thresholds between earnings bins. These cutoffs are chosen to be equally spaced on a log scale, with no interval less than \$5,000. Representative earnings are the average earnings in that bin, except for the lowest earnings bin, in which households are assumed to have zero earnings. For adults with positive earnings, the CPS earnings distribution is smoothed using an Epanechnikov kernel with a bandwidth of \$10,000.

assumption that the government wants to hold the total revenue (taxes net transfers) it currently collects from single adult households constant.

4 The Impact of Children's Utility on Optimal Taxes

We begin by examining the effect of including children's utility in the social planner's problem. Figure 1 shows how the optimal tax schedule changes when children's consumption is part of the planner's objective function. The left panel presents the optimal schedule assuming p = 0, i.e. that children have no resource needs. This is equivalent to a standard model that considers adults' utility from consumption but does not consider children's utility from consumption. The figure on the right presents the optimal schedule under our baseline specification, where p = 0.5. In this case, children consume half as many resources as adults, and the social planner considers the utility generated from both adults' consumption and children's consumption.









Notes: The y-axis denotes pre-tax earnings in thousands of dollars. The x-axis denotes post-tax income, calculated as the pre-tax earnings plus the optimal tax or transfer amount. For readability, the x-axis is cut off at \$200,000. Appendix Figure A3 presents the same chart with an extended x-axis. The gray dashed line represents the 45 degree line, with observations falling above this line denoting pre-tax earnings levels with net transfers and observations below denoting earnings levels with net taxes. The parameters are $\theta = 1$, $\eta_{parent} = 0.5$, $\eta_{childless} = 0.3$, $\epsilon = 0.33$, F = 0.7. In the left figure excluding children's utility, p = 0. In the right figure with children's utility, p = 0.5.

When the social planner only considers the utility of the adult, the optimal tax schedules for childless adults and single parents are very similar. The optimal demogrant is roughly \$17,000 for both family types (\$17,800 for childless adults and \$17,100 for single parent families), with net transfers until pre-tax earnings reach roughly \$30,000. Since the marginal utilities of income are identical across family types when p = 0, the difference in the optimal schedules here only reflects differences in the extensive margin labor supply elasticities between childless adults (0.3) and adults with children (0.5) and differences in the underlying initial earnings distributions.

Accounting for children's utility from consumption substantially changes the optimal tax

system for single parents. Under this schedule, single parents (with an assumed 2 children) would receive a demogrant of \$29,500 and receive net transfers until pre-tax earnings reach roughly \$70,000. The demogrant for childless adults only decreases slightly relative to the case with identical marginal utilities, falling to \$13,900. The planner is able to fund this large increase in transfers to single parents without a large decrease in transfers to childless adults because there are three times as many childless adults as single parents (Table 1).

Figure 2 shows the marginal tax rates on moving from one income level to the next under each model. Among single parent households, the marginal tax rate on moving from zero earnings to the first earnings level is distinctly lower than the marginal tax rate on moving from the first earnings level to the second, to encourage labor force participation (Jacquet, Lehmann, and Van der Linden 2013). This pattern is not as pronounced for childless adults, largely because we assume childless adults have a lower extensive margin labor supply elasticity. Marginal tax rates do not fall below zero for either group. In the model accounting for children's utility from consumption, shown on the right panel, marginal rates for childless adults are consistently higher than those for single parents at the same income level. This indicates that there is some net redistribution from childless adults to single parents.

5 Key parameters determining optimal transfers to families

We now turn to examining the impact of parameter assumptions on the optimal system. We first consider how assumptions about redistributive preferences (θ), economies of scale (F), and children's resource needs (p) affect optimal transfers, individually and in combination. We then explore the sensitivity of the model to assumptions about labor supply responses, which have long been recognized as a key parameter for determining optimal value of transfers. We focus on the effect of parameter choices on the demogrant for single



Figure 2: Marginal tax rates under optimal tax schedule

Notes: The x-axis denotes pre-tax earnings. The y-axis denotes the marginal tax rate on moving from each income level to the next income level above it under the optimal system. The marginal tax rates in this figure represent those under our baseline specification, with parameters $\theta = 1$, $\eta_{parent} = 0.5$, $\eta_{childless} = 0.3$, $\epsilon = 0.33$, F = 0.7. In the left figure excluding children's utility, p = 0. In the right figure with children's utility, p = 0.5.

parent households as a summary measure of the optimal system.

5.1 Importance of assumptions related to families' marginal utility of income

The optimal level of transfers to single parent households is heavily dependent on their marginal utility of income relative to childless adults. This is a function of children's resource needs (p), families' economies of scale (F), and the level of redistribution across income levels (θ) .

Figure 3 shows how the optimal demogrant for single parent households varies across a range of possible values for each of these three parameters. The parameter p ranges from 0 to 1, where p = 0 corresponds to a model that ignores children's consumption needs, p = 0.5 corresponds to the Census Bureau equivalence scale, and p = 1 corresponds to the Congressional Budget Office (CBO) equivalence scale. The parameter θ ranges from 0.25 to 2, where $\theta = 1$ corresponds to log utility, $\theta < 1$ corresponds to relatively lower levels

of redistribution, and $\theta > 1$ corresponds to relatively higher levels of redistribution. The parameter F ranges from 0.5 to 1, where F = 1 corresponds to families having no economies of scale in consumption, F = 0.7 corresponds to the economies of scale assumptions in the Census Bureau equivalence scale, and F = 0.5 corresponds to those in the CBO equivalence scale. The baseline specification assumes p = 0.5, $\theta = 1$, and F = 0.7.

Figure 3:	Sensitivity of a	optimal	demogrant f	for f	amilies t	o marginal	utilitv	parameters
()						()		

	p										
Θ	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1
F=0.5	<u> </u>										
0.25	8.3	14.2	20.7	27.0							
0.5	12.8	16.6	20.4	24.1	27.6	30.8	33.9				
0.75	15.4	18.6	21.6	24.4	27.1	29.6	31.9	34.1	36.1		
1	17.1	20.0	22.6	25.1	27.4	29.5	31.5	33.3	35.1	36.7	38.2
1.25	18.4	21.0	23.4	25.7	27.8	29.7	31.5	33.1	34.7	36.2	37.6
1.5	19.3	21.8	24.1	26.2	28.1	29.9	31.6	33.1	34.6	36.0	37.2
1.75	20.0	22.4	24.6	26.6	28.5	30.2	31.7	33.2	34.6	35.9	37.1
2	20.6	22.9	25.1	27.0	28.7	30.4	31.9	33.3	34.6	35.8	37.0
F=0.7											
0.25	8.3	12.6	17.1	21.6	25.9	29.7					
0.5	12.8	16.1	19.3	22.3	25.2	28.0	30.5	32.9			
0.75	15.4	18.4	21.2	23.8	26.3	28.6	30.8	32.8	34.7	36.5	
1	17.1	20.0	22.6	25.1	27.4	29.5	31.5	33.3	35.1	36.7	38.2
1.25	18.4	21.2	23.7	26.0	28.2	30.3	32.1	33.9	35.6	37.1	38.5
1.5	19.3	22.0	24.5	26.8	28.9	30.9	32.7	34.4	36.0	37.5	38.9
1.75	20.0	22.7	25.2	27.4	29.5	31.4	33.2	34.8	36.4	37.9	39.2
2	20.6	23.3	25.7	27.9	30.0	31.8	33.6	35.2	36.7	38.2	39.5
F=1											
0.25	8.3	10.4	12.4	14.4	16.3	18.0	19.7	21.3	22.8	24.2	25.5
0.5	12.8	15.3	17.6	19.8	21.9	23.9	25.7	27.5	29.1	30.6	32.1
0.75	15.4	18.1	20.6	23.0	25.2	27.2	29.2	31.0	32.7	34.3	35.7
1	17.1	20.0	22.6	25.1	27.4	29.5	31.5	33.3	35.1	36.7	38.2
1.25	18.4	21.3	24.1	26.6	28.9	31.1	33.1	35.0	36.8	38.5	40.0
1.5	19.3	22.3	25.2	27.7	30.1	32.4	34.4	36.4	38.1	39.8	41.4
1.75	20.0	23.1	26.0	28.6	31.1	33.3	35.4	37.4	39.2	40.9	42.5
2	20.6	23.8	26.7	29.3	31.8	34.1	36.2	38.2	40.0	41.7	43.3

Notes: Each cell represent the optimal transfer to a single parent households with two children and no earnings, in thousands of dollars, under varying parameter values. Darker colors denote larger transfers. All specifications use the same assumptions about the current tax system, the current earnings distribution, and labor supply responses as the baseline specification ($\eta_{parent} = 0.5$, $\eta_{childless} = 0.3$, $\epsilon = 0.33$). The baseline specification corresponds to p = 0.5, F = 0.7, $\theta = 1$.

Figure 3 shows that the value of the optimal demogrant for single parent households is especially sensitive to p. For example, under the baseline specification, moving from the Census Bureau's assumption that p = 0.5 to CBO's assumption that p = 1 increases the optimal demogrant for single parent households from \$29,500 to \$38,200 (30%). By contrast, doubling θ to 2 only increases the optimal demogrant for single parent households to \$31,800 (8%). Modifying F has no impact on the optimal demogrant under the baseline specification. Figure 3 confirms that these general patterns hold for parameter values outside the baseline specification – varying p across a range of reasonable values has an even larger impact on the optimal demogrant than varying θ or F across a range of reasonable values. This suggests that accurately accounting for children's consumption needs is of first order importance to optimal policy design and analysis.

The optimal demogrant for single parent households does not vary substantially with θ because the social planner has both equity and efficiency motives to transfer resources to low-income single parent households.¹⁵ The equity motive is driven by the fact that low-income single parent households have minimal resources and multiple people consuming them, so each person consumes very little. The efficiency motive is driven by the fact that single parent households generate more consumption per dollar of transfers than childless adults, due to economies of scale in consumption. When θ is high, equity motivates the social planner to transfer resources to low-income households, especially those with children. When θ is low, efficiency motivates the social planner to transfer resources from childless adults to single parent households. In both cases, the social planner wants to transfer resources to low-income single parents. Note this pattern does not hold if families do not have economies of scale in consumption, as represented by F = 1, because there is no longer an efficiency motive to transfer resources from childless adults to single parent.

5.2 Importance of assumptions related to single parents' labor supply responses

The magnitude of single parents' extensive margin labor supply responses is a focus of policy debates around the optimal transfers to single parents without earnings (Bastian 2023).

¹⁵There are some notable nonlinearities at low values of θ , where decreasing θ sharply increases the optimal demogrant. This is because the efficiency motive to transfer resources from childless households to families becomes very strong.

Figure 4 shows the sensitivity of the optimal demogrant amount to alternative assumptions about single parents' extensive margin labor supply responses. The baseline specification assumes single parents' extensive margin labor supply elasticity is 0.5, and that anyone with earnings up to \$30,000 can exit the labor force in response to tax changes. Reducing the single parent labor supply elasticity to 0.3, which is close to the lower end of estimates in the literature, increases the optimal demogrant for single parents' labor supply elasticity to 1 only reduces the optimal demogrant for single parent households to \$26,800 (a 9.2% decrease).

The optimal demogrant is less sensitive to the choice of single parent labor supply elasticity than to the choice of equivalence scale. For comparison, moving from the Census equivalence scale to the CBO equivalence scale increases the optimal demogrant for single parent households to \$38,200, or 29.5%.

6 Evaluation of current tax and transfer policy

We use the model to evaluate whether the current tax and transfer system allocates resources to single parent households in a manner that is consistent with assumptions that equivalence scales make regarding children's resource needs and family economies of scale. We summarize the current tax schedule from the relationship between earnings and taxes/transfers among single adult households without and without children in the CPS. We find that current policy is less generous to low-income households and households with children than our baseline specification of the optimal system. We then estimate the redistributive preferences that rationalize the current tax system as optimal. The current tax system places relatively little value on redistribution from higher-income households to lower-income households and from households without children to households with children. The level of redistribution from childless households to households with children is inconsistent with equivalance scales at Figure 4: Sensitivity of optimal demogrant to single parents' extensive margin labor supply elasticity



Notes: The x-axis denotes the value of single parents' extensive margin labor supply elasticity. The y-axis denotes the optimal transfer to households without earnings. All other model parameters follow the same values as the baseline specification (F = 0.7, p = 0.5, $\theta = 1$, $\eta_{childless} = 0.3$, $\epsilon = 0.33$).

most earnings levels, with the exception of low-earning working households.

6.1 Comparison of optimal policy to current policy

We compare our estimates of the optimal transfer system to estimates of cash and near-cash transfers that families can expect to receive under the current transfer system. We estimate average taxes and transfers by earnings level and family size using Supplemental Poverty Measure (SPM) units with only one adult in the March 2018-2020 CPS.¹⁶ Transfer income is

¹⁶Since transfer income is measured by SPM unit, it can only accurately be measured for childless adults and single parents that live alone. Adults that live with other adults (i.e. extended family or unmarried partners) are included in the optimal tax model, but not in these estimates of average taxes/transfers under current law. The sample is restricted to adults age 18-60 that are not students. Pre-tax earnings are measured from wages, farm, and self-employment income of the adult.

estimated based on the transfers used in calculating the supplemental poverty measure, which includes income from welfare/public assistance, SNAP, school lunch programs, WIC, housing subsidies, energy subsidies, SSI, the EITC, and disability and survivor's benefits, minus net state and federal taxes (Short 2011). Within each household type (childless adults and single parents), we estimate the current tax schedule by fitting a kernel smoothed function to the relationship between earnings and total taxes/transfers.¹⁷

Current policy offers lower transfers to households with children and households without earnings than is optimal, even under conservative assumptions. Figure 5 compares the estimates of current taxes/transfers to the optimal tax and transfer schedule from our baseline simulation. A non-earning childless adult can expect to receive \$4,300 transfers under the current system, which is well below the optimal demogrant of \$13,900. A non-earning single parent family can expect to receive \$9,700 in transfers, which is well below the optimal demogrant of \$29,500. Across all income groups, the optimal system generates a larger wedge between childless adults' tax liabilities and single parents' tax liabilities than the current system. Unlike the current system, the optimal system does not impose negative marginal tax rates on single parents with no earnings. This suggests that current policy underemphasizes childrens' consumption needs and overemphasizes the benefits of using tax and transfer policy to induce single parents to work.

There are two main limitations of using the CPS to estimate the current transfer system. The first is concerns about underreporting of benefits and income on surveys (Meyer, Mok, and Sullivan 2015; Meyer and Mittag 2019), which would bias our estimates of the generosity of the current system downwards. The second concern is with identifying the non-earning population. While wage income is a reasonable proxy for individuals with no income, some of the non-wage earners could have income or wealth from other sources that makes them ineligible for cash and cash-like transfers. Including these individuals in the estimate of

 $^{^{17}}$ The smoothed tax schedule is then discretized by evaluating expected taxes/transfers at the representative earnings level for each earnings bin from table 1. We choose a quadratic kernel and set the bandwidth to the lesser of \$10,000 or 20% of earnings. It is necessary to use a wider smoothing bandwidth at higher earnings levels because CPS data is more sparse.

Figure 5: Comparison of optimal tax schedule with current policy



Notes: The x-axis denotes pre-tax earnings. For readability, the x-axis is cut off at \$200,000. The y-axis denotes post-tax income, calculated as the pre-tax earnings plus the tax or transfer amount. The gray dashed line represents the 45 degree line, with observations falling above this line denoting pre-tax earnings levels with net transfers and observations below denoting earnings levels with net taxes. Results reflect the model with baseline specifications of F = 0.7, p = 0.5, $\theta = 1$, $\eta_{parents} = 0.5$, $\eta_{childless} = 0.3$, $\epsilon = 0.33$. Estimates of the current tax and transfer system come from supplemental poverty measure units with a single adult age 18-60 (not in school) in the pooled 2018-2020 March CPS. Pre-tax earnings are total wage, farm, and self-employment income. The tax and transfer amounts include reported income from welfare/public assistance, SNAP, school lunch programs, WIC, housing subsidies, energy subsidies, SSI, the EITC, and disability and survivor's benefits, minus net state and federal taxes.

the transfer to non-earners will also bias our estimate of current transfers to low-income households downwards. As an alternative we estimate current transfers to non-earners using statutory benefit amounts and estimates of take-up from the literature. This approach also concludes that current transfers to non-earning families are well below the optimal amount (Appendix B).

6.2 Redistributive preferences implied by current policy

To decompose how current policy diverges from our estimates of optimal policy, we invert the model to calculate the marginal social welfare weights $g_{Z,K}$ that rationalize the current tax and transfer schedule. These weights can be interpreted as the dollar value the government is willing to pay to increase each group's income by a dollar (Saez and Stantcheva 2016). Interpreting these results in a welfarist framework allows us to infer the government's beliefs about how the marginal utility of income varies by income and the presence of children, because a household's marginal social welfare weight is directly proportional to its marginal utility of income.

Figure 6 plots marginal social welfare weights by income and presence of children. In general, the welfare weights implied by the current system are larger for lower-income house-holds than for higher-income ones. For example, among childless households, the welfare weight of households in the top income group is 0.6, compared to 1.2 for households in the bottom income group. This indicates that, on the margin, the government values increasing the incomes of the lowest-income households twice as much as increasing the incomes of the highest-income households. Interpreted in a welfarist framework, this implies that the marginal utility of income for the lowest-income households is twice that of the highest-income households.

We can summarize the overall level of redistribution from high-income to low-income households by applying structure to the functional form of marginal utility. Under the CRRA utility function assumed in the optimal tax model, the relationship between after-tax income $I_{Z,K=0}$ and welfare weight $g_{Z,K=0}$ among childless households is a function of the parameter θ :

$$\lambda g_{Z,K=0} = M U I_{Z,K=0} = I_{Z,K=0}^{-\theta}$$
(23)

Fitting a logarithmic relationship between $I_{Z,K=0}$ and $g_{Z,K=0}$ under the current tax system



Figure 6: Marginal social welfare weights under current policy

Notes: The x-axis represents the after-tax income in each earnings group under the current tax system. The y-axis represents the marginal social welfare weight of each group, calculated from equation 12. The current earnings distribution is described in section 3.1.4 and the current tax system is described in section 6.1. Labor supply response assumptions match the baseline model specification ($\eta_{parent} = 0.5$, $\eta_{childless} = 0.3$, $\epsilon = 0.33$, extensive margin responses occur in income groups earning up to \$30,000).

suggests that the level of redistribution across income groups can best be rationalized by $\theta = 0.16$. This is a low level of redistribution relative to the conventional benchmark of $\theta = 1$, but is consistent with literature suggesting that the redistributive preferences that rationalize the United States tax system are low relative to past U.S. policy and low relative to other countries (Bargain et al. 2014; Lockwood and Weinzierl 2016; Hendren 2020).

We infer the current level of redistribution from childless households to households with children from the relationship between the welfare weights of single parent and childless households within the same income group. Figure 6 shows that under the current tax system, at the bottom of the income distribution, single parent and childless households have roughly equal welfare weights. By contrast, among low-income working households (those with after-tax incomes around \$30,000), single parent households are weighted almost twice as much as childless households. Among moderate-income households (those with after-tax incomes around \$40,000-\$50,000), single parent households are weighted well below childless households. Among higher-income households (those with after-tax incomes above roughly \$70,000), single parent and childless households are again weighted roughly equally.

These patterns are not consistent with the assumptions about children's resource needs and family economies of scale underlying the Census Bureau and CBO equivalence scales. Under the model assumptions, the social planner weights single parent households higher than childless households at every income level because single parent households have a higher marginal utility of income. Applying the functional form of marginal utility introduced in section 3.1.2, the ratio between single parent and childless welfare weights at a given income level I is a function of the equivalence scale parameters p and F, as well as the CRRA utility parameter θ :

$$\frac{g_{I,K=2}}{g_{I,K=0}} = \frac{MUI_{I,K=2}}{MUI_{I,K=0}} = (1+2p)^{1-(1-\theta)F}$$
(24)

To match the overall level of redistribution across income groups observed under the current tax system, we apply $\theta = 0.16$. Under the Census Bureau equivalence scale (p = 0.5, F = 0.7), single parent household welfare weights would be 32 percent higher than childless households at the same income level. Under the CBO equivalence scale (p = 1, F = 0.5), single parent household welfare weights would be 88 percent higher than childless households at the same income level. These are inconsistent with the current tax system, which weights single parent households equal to or even below childless households at most income levels. The exception is low-income working households (with after-tax incomes around \$30,000). In this group, single parent household welfare weights are almost twice that of childless households, which is in the range implied by the Census Bureau and CBO equivalence scales.

7 Discussion

This model motivates redistribution toward families with children based on the utility generated from children's consumption, but redistribution toward families could also be motivated by positive externalities from investment in children. A growing body of literature finds that transfers to low-income children can increase educational achievement, health, and earnings, which has positive effects on the government's budget (Hoynes and Patel 2018; Brown, Kowalski, and Lurie 2020; Bastian and Jones 2021). Certain programs pay for themselves on the margin – in many cases, an additional dollar of program spending on children generates more than an additional dollar of government revenue over the long term (Hendren and Sprung-Keyser 2020). Taking these positive externalities into account would increase the optimal amount of redistribution to low-income families with children.

The model does not explicitly include childcare costs, which may impact parents' labor supply responses to tax changes or working families' consumption levels. The simplest way to model childcare costs is as part of the disutility of work. In this case, the impact of childcare on labor supply decisions is captured in parents' labor supply elasticities, and childcare would have no impact on the marginal utility of income. Assuming existing empirical estimates capture parents' labor supply responses inclusive of childcare concerns, then our model sufficiently captures childcare by allowing single parents to have different labor supply elasticities than childless adults.

Another way childcare costs could enter the model is through a monetary cost in the family's budget constraint. Childcare spending could be modeled as a component of children's consumption, an investment in children's future ability, or simply a cost of work (Ho and Pavoni 2020; Bastani, Blomquist, and Micheletto 2020). If childcare spending is a component of children's consumption, it may be captured in our assumptions about how families allocate expenditures between adults and children. However, to the extent that non-working parents engage in home production of childcare, we may underestimate the resources of non-working families relative to working families. The model does not include married households, so all redistribution to single parents comes from single childless adults. The model holds total revenue raised from unmarried adults constant and abstracts away from redistribution between single and married households. We focus on unmarried households because the equivalence scales used by government agencies abstract away from any bargaining that occurs between adults in the household. The model could be extended to include married couples by making strong assumptions about how joint labor supply and the allocation of consumption within the household respond to changes in the tax schedule.

8 Conclusion

Redistribution of resources to individuals who cannot provide for themselves is an important feature of the United States' transfer policy. The optimal taxation literature often ignores this motivation for redistribution when it comes to children, focusing instead on children's future earning potential or on the utility of the parents. We estimate optimal taxes and transfers when children's utility from consumption is directly included in the social welfare function, allowing for redistribution to families on the basis of children's resource needs. We derive children's and family utility based on assumptions underlying equivalence scales used across government agencies, which allows us to model the optimal tax and transfer system under a set of transparent assumptions that are consistent with those used by the government in policy-making and poverty measurement.

Our findings demonstrate that accounting for children's consumption needs is of first order importance to analyzing the optimal tax and transfer system. We estimate that the optimal transfer to single parents without earnings is roughly \$29,500 when children's resource needs are included in the planner's problem, compared to just \$17,100 when those needs are not considered. The optimal tax and transfer system is even more sensitive to assumptions about children's resource needs than assumptions about the overall level of redistribution between income groups or the magnitude of labor supply distortions. Improving estimates of children's resource needs and economies of scale within families is an important task for future work, as these are key inputs to policymaking and optimal policy analysis.

References

- Aizer, Anna, Hilary Hoynes, and Adriana Lleras-Muney. 2022. "Children and the US Social Safety Net." The Journal of Economic Perspectives 36: 149–174.
- Apps, Patricia, and Ray Rees. 2018. "Optimal family taxation and income inequality." International Tax and Public Finance 25: 1093–1128.
- Baker, Michael, Derek Messacar, and Mark Stabile. 2023. "Effects of Child Tax Benefits on Poverty and Labor Supply: Evidence from the Canada Child Benefit and Universal Child Care Benefit." *Journal of Labor Economics* 41: 1129 – 1182.
- Balestrino, Alessandro, Alessandro Cigno, and Anna Pettini. 2002. "Endogenous Fertility and the Design of Family Taxation." International Tax and Public Finance 9: 175–193.
- Bargain, Olivier, Mathias Dolls, Dirk Neumann, Andreas Peichl, and Sebastian Seigloch. 2014. "Comparing Inequality Aversion Across Countries When Labor Supply Responses Differ." International Tax and Public Finance 21: 845–873.
- Bastani, Spencer, Soren Blomquist, and Luca Micheletto. 2020. "Child Care Subsidies, Quality, and Optimal Income Taxation." *American Economic Journal: Economic Policy* 12: 1–37.
- Bastian, Jacob. 2020. "The Rise of Working Mothers and the 1975 Earned Income Tax Credit." *American Economic Journal: Economic Policy* 12: 44–75.
- Bastian, Jacob. 2023. "How Would a Permanent 2021 Child Tax Credit Expansion Affect Poverty and Employment?" National Tax Journal: Forthcoming, https://drive.google.com/file/d/ 122-yvQ_OCvNp6em-o8rsioO79TYQL3vp/view.
- Bastian, Jacob, and Maggie Jones. 2021. "Do EITC expansions pay for themselves? Effects on tax revenue and government transfers." *Journal of Public Economics* 196.
- Baughman, Reagan, and Stacy Dickert-Conlin. 2009. "The earned income tax credit and fertility." Journal of Population Economics 22: 537–563.
- Betson, David M. 1990. "Alternative Estimates of the Cost of Children from the 1980-86 Consumer Expenditure Survey."
- Betson, David M. 1996. ""Is Everything Relative?" The Role of Equivalence Scales in Poverty Measurement." https://www.census.gov/library/working-papers/1996/demo/betson-02.html.

- Blumkin, Tomer, Yoram Margalioth, and Efraim Sadka. 2015. "The Re-distributive Role of Child Benefits Revisited." International Tax and Public Finance 22: 476–501.
- Blundell, Richard, and Andrew Shephard. 2012. "Employment, Hours of Work and the Optimal Taxation of Low-Income Families." 79: 481–510.
- Bourguignon, Francois, and Amedeo Spadaro. 2012. "Tax Benefit Revealed Social Preferences." The Journal of Economic Inequality 10: 75–108.
- Brown, David W, Amanda E Kowalski, and Ithai Z Lurie. 2020. "Long-Term Impacts of Childhood Medicaid Expansions on Outcomes in Adulthood." The Review of Economic Studies 87: 792– 821.
- Browning, Martin, Pierre André Chiappori, and Arthur Lewbel. 2013. "Estimating consumption economies of scale, adult equivalence scales, and household bargaining power." Review of Economic Studies 80: 1267–1303.
- Calvi, Rossella, Jacob Penglase, Denni Tommasi, and Alexander Wolfe. 2023. "The more the poorer? Resource sharing and scale economies in large families." *Journal of Development Economics* 160.
- Casarico, Alessandra, Luca Micheletto, and Alessandro Sommacal. 2015. "Intergenerational transmission of skills during childhood and optimal public policy." *Journal of Population Economics* 28: 353–372.
- Chetty, Raj. 2012. "Bounds on Elasticities with Optimization Frictions: A Synthesis of Micro and Macro Evidence in Labor Supply." *Econometrica* 80: 969–1018.
- Chetty, Raj, John N Friedman, and Emmanuel Saez. 2013. "Using Differences in Knowledge Across Neighborhoods to Uncover the Impacts of the EITC on Earnings." *American Economic Review* 103: 2683–2721.
- Chiappori, Pierre-Andre, and Costas Meghir. 2015. "Intrahousehold Inequality." In Handbook of Income Distribution, vol. 2, 1369–1418.
- Chiappori, Pierre André. 2016. "Equivalence Versus Indifference Scales." *The Economic Journal* 126: 523–545.
- Choné, Philippe, and Guy Laroque. 2010. "Negative Marginal Tax Rates and Heterogeneity." American Economic Review 100: 2532–47.

- Cigno, Alessandro, Annalisa Luporini, and Anna Pettini. 2004. "Hidden information problems in the design of family allowances." *Journal of Population Economics* 17: 645–655.
- Cigno, Alessandro, Annaliza Luporini, and Anna Pettini. 2003. "Transfers to families with children as a principal-agent problem." *Journal of Public Economics* 87: 1165–1177.
- Cigno, Alessandro, Pierre Pestieau, and Ray Rees. 2011. "Introduction to the Symposium on Taxation and the Family." *CESifo Economic Studies* 57: 203–215.
- Cigno, Alessandro, and Anna Pettini. 2002. "Taxing family size and subsidizing child-specific commodities?" Journal of Public Economics 84: 75–90.
- Cremer, Helmuth, Arnaud Dellis, and Pierre Pestieau. 2003. "Family size and optimal income taxation." *Journal of Population Economics* 16: 37–54.
- Dunbar, Geoffrey R, Arthur Lewbel, and Krishna Pendakur. 2013. "Children's Resources in Collective Households: Identification, Estimation, and an Application to Child Poverty in Malawi." *American Economic Review* 103: 438–471.
- Edin, Kathryn, and Maria Kefalas. 2005. Promises I can keep: Why poor women put motherhood before marriage.: University of California Press.
- Eissa, Nada, and Jeffrey B Liebman. 1996. "Labor Supply Response to the Earned Income Tax Credit." The Quarterly Journal of Economics 111: 605–637.
- Feinberg, Robert, and Daniel Kuehn. 2018. "Guaranteed Nonlabor Income and Labor Supply: The Effect of the Alaska Permanent Fund Dividend." The B.E. Journal of Economic Analysis Policy 18: 13.
- Gelber, Alexander, and Matthew Weinzierl. 2016. "Optimal Taxation When Children's Abilities Depend on Parents' Resources." National Tax Journal 69: 11–39.
- Giannarelli, Linda. 2019. "What Was the TANF Participation Rate In 2016? How Has the Participation Rate Changed over Time?." https://www.urban.org/research/publication/ what-was-tanf-participation-rate-2016.
- Guner, Nezih, Remzi Kaygusuz, and Gustavo Ventura. 2020. "Child-Related Transfers, Household Labour Supply, and Welfare." The Review of Economic Studies 87: 2290–2321.
- Hendren, Nathaniel. 2020. "Measuring Economic Efficiency Using Inverse-Optimum Weights." Journal of Public Economics 187.

- Hendren, Nathaniel, and Ben Sprung-Keyser. 2020. "A Unified Welfare Analysis of Government Policies." The Quarterly Journal of Economics 135: 1209–1318.
- Ho, Christine, and Nicola Pavoni. 2020. "Efficient Child Care Subsidies." American Economic Review 110: 162–99.
- Horvath-Rose, Ann E, H Elizabeth Peters, and Joseph J Sabia. 2008. "Capping Kids: The Family Cap and Nonmarital Childbearing." Source: Population Research and Policy Review 27: 119– 138, https://www.jstor.org/stable/41217941.
- Hoynes, Hilary, and Ankur Patel. 2018. "Effective Policy for Reducing Poverty and Inequality? The Earned Income Tax Credit and the Distribution of Income." *Journal of Human Resources* 53: 44.
- Jacquet, Laurence, Etienne Lehmann, and Bruno Van der Linden. 2013. "Optimal redistributive taxation with both extensive and intensive responses." Journal of Economic Theory 148: 1770– 1805.
- Joint Center for Housing Studies of Harvard University. 2020. "America's Rental Housing 2020." https://www.jchs.harvard.edu/americas-rental-housing-2020.
- Jones, Damon, and Ioana Marinescu. 2022. "The Labor Market Impacts of Universal and Permanent Cash Transfers: Evidence from the Alaska Permanent Fund." American Economic Journal: Policy 14: 315–40.
- Jones, Maggie R., and Amy B. O'Hara. 2017. "Do doubled-up families minimize household-level tax burden?" *National Tax Journal* 69: 613–640.
- Joyce, Ted, Robert Kaestner, Sanders Korenman, and Stanley Henshaw. 2004. "Family Cap Provisions and Changes in Births and Abortions." *Population Research and Policy Review* 23: 475–511.
- Kearney, Melissa Schettini. 2004. "Is There an Effect of Incremental Welfare Benefits on Fertility Behavior? A Look at the Family Cap." Journal of Human Resources 56: 295–325.
- Kurnaz, Musab. 2021. "Optimal Taxation of Families: Mirrlees Meets Becker." The Economic Journal 131: 2984.
- Laroque, Guy. 2005. "Income Maintenance and Labor Force Participation." *Econometrica* 73: 341–376.

- Lewbel, Arthur, and Krishna Pendakur. 2008. "Estimation of collective household models with Engel curves." *Journal of Econometrics* 147: 350–358.
- Lippold, Kye, and Beata Luczywek. 2023. "Estimating Income Effects on Earnings using the 2021 Child Tax Credit Expansions." https://beata-luczywek.com/files/Luczywek_JMP.pdf.
- Lockwood, Benjamin B, and Matthew Weinzierl. 2016. "Positive and normative judgments implicit in U.S. tax policy, and the costs of unequal growth and recessions." Journal of Monetary Economics 77: 30–47.
- Menon, Martina, Krishna Pendakur, and Federico Perali. 2012. "On the expenditure-dependence of children's resource shares." *Economics Letters* 117.
- Meyer, Bruce D, and Nikolas Mittag. 2019. "Using linked survey and administrative data to better measure income: Implications for poverty, program effectiveness, and holes in the safety net." *American Economic Journal: Applied Economics* 11: 176–204.
- Meyer, Bruce D, Wallace K C Mok, and James X Sullivan. 2015. "The Under-Reporting of Transfers in Household Surveys: Its Nature and Consequences."
- Meyer, Bruce D, and Dan T Rosenbaum. 2001. "Welfare, the Earned Income Tax Credit, and the Labor Supply of Single Mothers." The Quarterly Journal of Economics 116: 1063–1114, https://www.jstor.org/stable/2696426.
- Miller, Cynthia, Lawrence F Katz, Gilda Azurdia, Adam Isen, Caroline Schultz, and Kali Aloisi. 2018. "Boosting the Earned Income Tax Credit for Singles: Final Impact Findings from the Paycheck Plus Demonstration in New York City." https://www.mdrc.org/work/publications/ boosting-earned-income-tax-credit-singles.
- Moffitt, Robert. 2015. "The U.S. Safety Net and Work Incentives: The Great Recession and Beyond." Journal of Policy Analysis and Management 34: 458–466, https://www.jstor.org/stable/ 43866380.
- Parolin, Zachary. 2021. "Decomposing the Decline of Cash Assistance in the United States, 1993 to 2016." Demography 58: 1119–1141.
- Saez, Emmanuel. 2002. "Optimal Income Transfer Programs: Intensive versus Extensive Labor Supply Responses." The Quarterly Journal of Economics 117: 1039–1073.
- Sauval, Maria, Greg J Duncan, Lisa A Gennetian, Katherine A Magnuson, Nathan A Fox, Kim-

berly G Noble, and Hirokazu Yoshikawa. 2024. "Unconditional cash transfers and maternal employment: Evidence from the Baby's First Years study." *Journal of Public Economics* 236.

- Schwabish, Jonathan. 2007. "Take a penny, leave a penny: The propensity to round earnings in survey data." Journal of Economic and Social Measurement 32: 93–111.
- Short, Kathleen. 2011. "Supplemental Poverty Measure: 2010." https://www.census.gov/library/publications/2011/demo/p60-241.html.
- Shrider, Em. 2021. "Alternative School Lunch Valuation in the CPS ASEC During COVID-19." https://www.fns.usda.gov/nslp/nslp-fact-sheet.
- Splinter, David, Jeff Larrimore, and Jacob Mortenson. 2017. "Whose child is this? Shifting of dependents among EITC claimants within the same household." National Tax Journal 70: 737–758.
- USDA. 2017. "SNAP Fiscal Year 2018 Cost-of-Living Adjustments."
- Vivalt, Eva, Elizabeth Rhodes, Alexander W Bartik, David E Broockman, and Sarah Miller. 2024.
 "The Employment Effects of a Guaranteed Income: Experimental Evidence from Two U.S. States."

Appendix

N

12

Share of households

02

c

ò

A Appendix Figures

100



0

ò

•

400



Childless adults

200

Pre-tax earnings (thousands \$)

300

Single parents

200 Pre-tax earnings (thousands \$)

300

100

+ (

400

Notes: The x-axis denotes the representative pre-tax earnings in each earnings bin in thousands of dollars. The y-axis denotes the share of households in each earnings bin. Earnings distribution estimates are from pooling the 2018-2020 March CPS. The sample is unmarried adults, restricted to those age 18-60 and excluding students. Workers with negative earnings or with investment income greater than a quarter of earned income are excluded from the sample. The "before smoothing" distributions represent the share of households in each earnings bin and the average earnings in each bin directly from the CPS. To generate the "smoothed" distributions, among adults with non-zero earnings, the CPS earnings distribution is smoothed using an Epanechnikov kernel with a bandwidth of \$10,000.



Figure A2: Robustness to more detailed approximation of current tax system

Notes: The specification shown in this figure initializes the model using the detailed summary of the current tax system shown in Section 6.1, rather than the linear approximation of the current tax system used in the baseline specification. The x-axis denotes pre-tax earnings and is cut off at \$200,000 for readability. The y-axis denotes post-tax income, calculated as the pre-tax earnings plus the optimal tax or transfer amount. The gray dashed line represents the 45 degree line, with observations falling above this line denoting pre-tax earnings levels with net transfers and observations below denoting earnings levels with net taxes. Otherwise, this model matches the baseline specification, with parameters $\theta = 1$, F = 0.7, p = 0.5, $\eta_{parents} = .5$, $\eta_{childless} = .3$, and $\epsilon = .33$. The demogrant is \$13,900 for childless adults and \$30,400 for single parent households.



Figure A3: Optimal tax schedule to \$400,000 pre-tax earnings

Excluding children's utility

With children's utility

Notes: The y-axis denotes pre-tax earnings in thousands of dollars. The x-axis denotes post-tax income, calculated as the pre-tax earnings plus the optimal tax or transfer amount. The gray dashed line represents the 45 degree line, with observations falling above this line denoting pre-tax earnings levels with net transfers and observations below denoting earnings levels with net taxes. The parameters are $\theta = 1$, $\eta_{parent} = 0.5$, $\eta_{childless} = 0.3$, $\epsilon = 0.33$, F = 0.7. In the left figure excluding children's utility, p = 0. In the right figure with children's utility, p = 0.5.





Notes: The x-axis denotes the maximum pre-tax earnings level assumed to have a positive extensive margin elasticity. The y-axis denotes the optimal transfer to non-earning households. The black vertical dashed line denotes the baseline value of extensive margin elasticities up to \$30,000 in pre-tax earnings. The remaining parameters match the baseline specification (p = 0.5, F = 0.7, $\theta = 1$, $\eta_{parent} = 0.5$, $\eta_{childless} = 0.3$, $\epsilon = 0.33$).

B Statutory Benefits to Non-earners

Given the measurement issues in the CPS, we also evaluate the current system compared to the optimal system using statutory benefit amounts and focusing on the value of the demogrant. However, given that two of the main programs for families without earnings (TANF and housing assistance) have limited funding, statutory benefit amounts alone are not a good proxy for the actual benefits a non-earning family can expect to receive. Federal TANF funding is fixed at 1996 nominal levels, and funding for housing assistance is capped by Congress annually. Both TANF and housing assistance only serve roughly a quarter of eligible households, and states are given flexibility in how they allocate these limited resources (Giannarelli 2019; Joint Center for Housing Studies of Harvard University 2020). Instead, this approach estimates the value of the transfers non-earning families receive *in expectation* using statutory benefit amounts and estimates from the literature on coverage rates, which reflect that the full benefits are not available to all eligible individuals. We focus on the four main programs that determine eligibility primarily based on income: TANF, SNAP, school lunch subsidies, and housing subsidies.

We assume households can expect to receive 100% of the *entitlement* program benefits they are income-eligible for (SNAP and school lunch), but scale *non-entitlement* program benefits (TANF and housing assistance) by estimates of the share of eligible households that receive benefits. Low levels of program receipt among eligible households reflect some combination of household decisions to take up programs and policy decisions to ration benefits. These are not cleanly separable, since states may reduce take-up by increasing the costs associated with getting and keeping benefits, such as frequent redetermination periods, lifetime limits, or long wait times (see Parolin (2021) for more examples in the TANF setting). Therefore, the share of households that receive benefits is the best available proxy for the average benefits a household could expect to receive.

Table B1 shows what these statutory benefit amounts and benefit coverage rates imply for the current value of the demogrant. If families were able to receive the full statutory benefit amount for these programs, the demogrant for childless adults would be \$9,738 and the demogrant for a single parent with two children would be \$26,013, which are broadly similar to the optimal demogrants in our baseline specification. However, after adjusting for incomplete coverage, the expected demogrant falls to \$2,435 for childless adults and \$11,531 for single parent families, well below the optimal values. These values are very similar to the demogrant amount estimated in the CPS.

	Childless adult	Single parent with 2 kids
Annual SNAP	NA	6,048
SNAP coverage rate	-	1
Annual TANF	NA	$5,\!688$
TANF coverage rate	-	0.25
Annual school lunch	NA	663
School lunch coverage rate	-	1
Annual housing	9,738	13,614
Housing coverage rate	0.25	0.25
Expected demogrant	2,435	11,531

Table B1: Expected demogrant from statutory benefit amounts

Notes: Maximum SNAP monthly benefits were \$192 for a single, childless adult and \$504 for a 3 person household in 2018 (USDA 2017). However, SNAP has work requirements for able-bodied adults without dependents, so a single childless adult that is not working and is not categorically eligible (due to disability or age) is not eligible for SNAP benefits. Maximum TANF benefits for a 3-person household vary by state. In our calculations we take the average maximum monthly benefit across all states as of July 2019, which is \$474. The data come from Tables II.A.4 and II.A.5 of the July 2019 Welfare Rules Databook. The annual value of school lunch is Census' estimate of the value of a full school year of lunches in 2020 (Shrider 2021). We calculate the value of housing benefits by averaging HUD data on fair market rents in October 2018, weighted by 2010 county population. We use studio rents for childless adults and 2-bedroom rents for single parents. The TANF coverage rate comes from Giannarelli (2019), and the housing coverage rate comes from Joint Center for Housing Studies of Harvard University (2020).