Econometrica, Vol. 91, No. 6 (November, 2023), 2261–2294

RETHINKING THE WELFARE STATE

NEZIH GUNER
CEMFI

REMZI KAYGUSUZ
Department of Economics, Durham University and Faculty of Arts and Social Sciences, Sabanci University

GUSTAVO VENTURA
Department of Economics, Arizona State University

The United States spends significant amounts on non-medical transfers for its working-age population in a wide range of programs that support low- and middle-income households. How valuable are these programs for U.S. households? Are there simpler, welfare-improving ways to transfer resources that are supported by a majority? What are the macroeconomic effects of such alternatives? We answer these questions in an equilibrium, life-cycle model with single and married households who face idiosyncratic productivity risk, in the presence of costly children and potential skill losses of females associated with non-participation. Our findings show that a potential revenue-neutral elimination of the welfare state generates large welfare losses in the aggregate, although most households support the move as losses are concentrated among a small group. We find that a *Universal Basic Income* program does not improve upon the current system. If, instead, per-person transfers are implemented alongside a proportional tax, a *Negative Income Tax* experiment, it becomes feasible to improve upon the current system. Providing per-person transfers to all households is costly, and reducing tax distortions helps to provide for resources to expand redistribution.

KEYWORDS: Taxes and transfers, universal basic income, household labor supply, income risk, social insurance.

1. INTRODUCTION

IN THIS PAPER, we focus on the set of means-tested government transfers available to households of working age in the United States. These transfers are sizable and cover a wide range of programs and tax credit provisions. We refer to them as the *welfare state* for short. We ask: to what extent do households value the current welfare state in the U.S.? Are there simpler, welfare-improving ways to transfer resources that are supported by a majority? What are the macroeconomic effects of switching to such alternatives?

Several observations motivate our work. First, the welfare state is far from insignificant: excluding healthcare transfers (Medicaid), spending in all different programs add up to

Nezih Guner: nezih.guner@cemfi.es
Remzi Kaygusuz: remzi.kaygusuz@durham.ac.uk
Gustavo Ventura: gustavo.ventura@asu.edu

Guner acknowledges support from EU 7th Framework Collaborative Project Integrated Macro-Financial Modeling for Robust Policy Design (MACFINROBODS), Grant 612796, the Spanish Ministry of Economy and Competitiveness Grants ECO2014-54401-P. We thank seminar and workshop participants at Bank of Spain, BGSE Summer Forum (Income Dynamics and Family), Boğaziçi University, Collegio Carlo Alberto, EUI, Federal Reserve Banks of Atlanta and Minneapolis, Georgetown U., Helsinki School of Economics, IMF, LAEF-Santa Barbara, Macroeconomics of Inequality Conference (St. Louis Fed), NBER Summer Institute (Macro Public Finance), Queen Mary U., SED Annual Meetings, SITE (Macroeconomics and Inequality), and Universities of Barcelona, Bristol, Cambridge, Connecticut, Emory, Geneva, Helsinki, Lugano, Purdue, Rice, Texas A&M, and Southern California. Francesco Chiocchio, Eugenio Renedo, Juan Carlos Ruiz Garcia, and Micole De Vera provided excellent research assistance. The usual disclaimer applies.

 https://doi.org/10.3982/ECTA19921

nearly 2.5% of GDP.[1] The rules and details of various programs are routinely discussed as key in affecting labor supply, inequality, and well-being in different ways. Hence, reforms or expansions of the current scheme are expected to have significant aggregate, distributive, and welfare effects. Second, most households are potentially two-earner households.[2] This matters as current transfers depend critically on marital status/gender differences and the presence of children. Furthermore, households with two potential earners can cope with labor market shocks better than single-person households. As a result, social insurance and redistribution policy recommendations for an economy with two (potential) earners are likely to be different than those for a single-earner economy. Lastly, marital status and gender differences are usually not considered in the analysis of tax and transfer policies. In particular, differences by marital status and gender in wage and earnings inequality over the life cycle are typically ignored. In this paper, we fill a void by providing a macroeconomic analysis that considers all these aspects. We do so by developing an equilibrium framework with uninsurable shocks, labor supply decisions in two-earner households, costly children, and a detailed representation of taxes and transfers.

We build an equilibrium life-cycle model with a number of novel features. First, we introduce a rich degree of heterogeneity in our model economy. Individuals differ by skill (i.e., education levels), gender, and marital status. Skilled and unskilled individuals face distinct wage rates and differ on how fast their skills evolve as they age. In addition, single and married individuals face permanent shocks at birth and uninsurable persistent shocks over their life cycle. Second, we allow for labor supply decisions of spouses at the extensive and intensive margins. Third, in line with data, we jointly account for the presence of children across married and single households, the timing of their arrival, and the associated childcare costs. In particular, we account for the level and variation of childcare costs over the life cycle as crucial determinants of female labor supply. Finally, we model the dynamic costs and benefits of participation decisions by allowing the labor market skills of females to depreciate due to childbearing disruptions.

Our parameterized model takes into account the different programs that comprise the U.S. welfare state and the progressive income tax system, excluding healthcare transfers, for example, Medicaid and Medicare. Transfers in the model economy consist of three main components. The first is the Earned Income Tax Credit that provides a refundable tax credit to households with earnings. The second component relates to child-related transfers, for example, the Child Tax Credit and childcare subsidies. The last part consists of the means-tested transfers, which are typically identified with the "welfare" system in the United States, for example, Temporary Assistance to Needy Families and Food Stamps. How much transfers households receive from different programs crucially depends on their marital status, earnings, number of children, and childcare expenses, and this dependence motivates our modeling choices. As such, a detailed description of the welfare state is a crucial input in the analysis. Any reform creates winners and losers, and the magnitude of these gains and losses critically depends on who benefits from the current system.

Given the welfare state and the tax system, we parameterize our model using U.S. aggregate and cross-sectional data. Our model economy is in line with how earnings inequality evolves over the life cycle (by gender, skill, and marital status), the levels and

[1]To place this number in international perspective, note that OECD (2019) calculates that income support to working-age population as a fraction of GDP was 1.9% in the U.S. The numbers for several European countries are much higher: Germany (3.5%), France (5.4%), Belgium (7.5%).

[2]More than 60% of the U.S. labor force between ages 25 and 54 is married (Current Population Survey, 2000–2018).

life-cycle changes in married females' participation rates, the life-cycle patterns of the gender wage-gap, and the rise in consumption dispersion with age. Altogether, our model economy presents a comprehensive macroeconomic model suitable to address the role and reforms of the welfare state.

Findings. We conduct three sets of experiments. First, we consider the hypothetical complete elimination of the welfare state and concomitantly reduce the income taxes for all households to achieve budget balance. This allows us to gauge the aggregate effects of the welfare state, and the valuation of the welfare state vis-a-vis a reduction in the tax burden. Overall, eliminating the welfare state leads to an increase in hours worked and participation rates of married females of about 3% and 4.5%, respectively, and an increase in output of about 1.7%. We find that eliminating the welfare state leads to a sharp aggregate welfare loss measured by a consumption compensating variation, of about 3.2% for a newborn individual under the veil of ignorance. Quite interestingly, a substantial majority of newborns support the hypothetical elimination of the welfare state (about 60.7%). This reflects the targeted nature of the current system, which is highly valuable to poor households and in particular to poor single mothers with children, while the majority of households either do not benefit from it or do so marginally.

We then introduce two major reforms to the welfare state. First, we replace the entire welfare state with a single transfer per person. We dub this case a *Universal Basic Income* (UBI). We search across steady states for the level of the transfer and the level of taxation that maximize ex ante welfare (under the veil of ignorance) that keeps the budget balanced. We find that a generous transfer per person of about 3.2% of mean household income (about \$3140 per person or \$12,550 for a family of four in 2019 dollars) maximizes the welfare of newborns.[3] However, even this welfare-maximizing level of transfers leads to an aggregate welfare *loss* of 1.3%. that is, there is no UBI program that can improve upon the current system. If we introduce a UBI scheme on top of the current welfare state, as most proponents of a UBI advocate, only small transfers lead to welfare gains. A relatively small transfer of just 1% of household income leads to welfare losses with a majority of individuals against such a program. Overall, our findings indicate that a UBI scheme is hardly a good idea in welfare terms.

Second, we replace all transfers and current income taxes with a single transfer per person and a proportional tax rate. We dub this case a *Negative Income Tax* (NIT). This case then combines a drastic transfer reform with a drastic tax reform. Similarly to the UBI case, we search across steady states for the level of the transfer and the associated tax rate that maximize the ex ante welfare of newborns and satisfy budget balance. We find that a generous transfer of about 4.8% of mean household income (about \$4700 per person or \$18,800 for a family of four in 2019 dollars) maximizes ex ante welfare with a gain of 0.2% and leads to strong majority support among newborns (about 68.2%). If a reform allows NIT transfers to differ between single and married households with more generous payments to singles, the welfare gains are larger (0.7%) and the program still has majority support of newborns (51.4%). The desirability of the NIT scheme becomes stronger when we take into account transitions across steady states.

The upshot for the relative success of an NIT scheme is that a larger degree of redistribution is feasible given the smaller tax distortions that ensue in this case. As tax distortions are reduced with a proportional tax, the size of the aggregate economy grows alongside the needed tax revenue to finance larger transfers. Therefore, an NIT scheme

[3]The mean household income in 2019 was about \$98,000.

makes higher degrees of redistribution feasible. We show that the desirability of an NIT scheme is resilient to variations in the environment.

In placing our findings in perspective, one important case that we consider is an NIT reform in an economy with a different underlying inequality. The U.S. economy has changed significantly during recent decades, with a sustained increase in inequality. Meanwhile, U.S. households changed with rising educational attainment, higher female labor force participation, and greater assortative mating. How do these changes affect the value of transfers for all and the desirability of an NIT reform? To answer this question, we recalibrate our model to salient features of the U.S. economy in the 1980s and introduce a NIT reform. We find that the NIT reform leads to even more significant welfare gains in such an environment. With less inequality, there is a lower demand for transfers and the welfare-maximizing NIT transfer becomes smaller. As a result, the required proportional tax rate is also smaller and larger welfare gains are possible. In the 1980 economy, 80% of newborn households favor the NIT scheme, while the support in the baseline is 68.2%. Thus, introducing a NIT scheme in an economy with today's characteristics is more difficult since the demand for transfers is greater, and financing such larger transfers requires larger tax rates.

Related Literature. Our paper is closely related to the literature that studies the welfare and aggregate effects of taxes and transfers in dynamic, general-equilibrium models with heterogeneous agents. Recent papers in this literature include Guner, Lopez-Daneri, and Ventura (2016), Heathcote, Storesletten, and Violante (2017), Badel, Huggett, and Luo (2020), Kindermann and Krueger (2020), Boar and Midrigan (2022), and Ferriere, Grubener, Navarro, and Vardishvili (2023). Within this literature, Ferriere et al. (2023) emphasized how larger transfers can be financed by lower progressivity of income taxes. Kaygusuz (2010, 2015), Guner, Kaygusuz, and Ventura (2012, 2020), Ortigueira and Siassi (2013), Holter, Krueger, and Stepanchuk (2019), Wu and Krueger (2021), and Borella, De Nardi, and Yang (2023), among others, considered environments with two-earner households. Ortigueira and Siassi (2022) studied how transfers can affect cohabitation versus marriage incentives, and Low, Meghir, Pistaferri, and Voena (2022) analyzed how marriage prospects can impact decisions to participate in programs with time limits. Blundell, Pistaferri, and Saporta-Eksten (2016) provided empirical evidence on the importance of family labor supply for consumption smoothing.

The UBI and its close cousin NIT have a long intellectual history (Moffitt (2003)) and gained support in recent public debate. Van Parijs and Vanderborght (2017) and Hoynes and Rothstein (2019) provided excellent reviews. Within macro-public-finance literature, Lopez-Daneri (2016) found that an NIT transfer of about 11% of mean income leads to a large, 2.1%, welfare gain despite sharp output losses. Luduvice (2021) and Conesa, Li, and Li (2023) considered replacing current transfers with a UBI and found that welfare gains are hard to achieve, as we find in this paper. Using search and matching models, Jaimovich, Saporta-Eksten, Setty, and Yedid-Levi (2022) also found that UBI implies welfare losses, while Rauh and Santos (2022) suggested that welfare gains are possible if UBI also replaces the current unemployment benefits. Daruich and Fernandez (2023) studied a UBI experiment within an overlapping generations model where the next generation's human capital depends on parents' decisions and found that UBI is not a good idea when the welfare of future generations is taken into account.

Our analysis differs from these papers on three key aspects. First, we provide novel facts on how inequality along the life cycle changes for individuals and households of different marital status and skill levels and use them to discipline the benchmark economy.

Second, the model economy features a comprehensive welfare state, necessary to identify winners and losers in any reform. Finally, the model economy consists of single and married households, and married females who make participation decisions. These features are critical to understanding the implications of any reform to the current transfer system since female labor supply responds significantly to changes in tax-transfer policies, and the current welfare system treats different households (married/single, with and without children) differently.

The paper is organized as follows. In Section 2, we document patterns of hours, earnings, and consumption over the life cycle of individuals and household in the United States. Section 3 presents the model economy. In Section 4, we describe the parameterization and calibration of the benchmark economy. Section 5 discusses the properties of the benchmark economy. In Section 6, we present the main findings of our quantitative experiments. Section 7 places our findings in perspective. Section 8 concludes. Throughout the paper, references are made to the online Supplemental Material (Guner, Kaygusuz, and Ventura (2023a)). This Supplemental Material also provides the replication packages for empirical and quantitative analysis in the paper.

2. EARNINGS, HOURS, AND CONSUMPTION: LIFE-CYCLE FACTS

We use the March Supplement of the CPS from 1980 to 2019 to document how average hourly wages, inequality of hourly wages and earnings, and labor market statistics (hours and participation) change over the life cycle. For age profiles for nondurable consumption, we use the Consumption Expenditure Survey (CEX) from 1984 to 2019. Our measure of inequality is the variance of logs. The Supplemental Material provides sample restrictions and the definitions of all the variables.

We estimate age profiles using repeated cross sections in the data. To this end, let $m_{j,t,c}$ be any statistic of interest for an age-j individual (or household) at time t, of cohort c. For example, $m_{j,t,c}$ could be the variance of log hourly wages among $j = 30$-year-olds in 2000, who are born in $c = t - j = 1970$, that is, the variance within a (j, t, c)-cell. Since age, time, and cohort are linearly dependent, we construct age profiles using two approaches. We first consider a time-effects specification by regressing $m_{j,t,c}$ on a set of age and time (year) dummy variables, that is,

$$m_{j,t,c} = \beta_j' \mathbf{D}_j + \beta_t' \mathbf{D}_t + \varepsilon_{j,t,c}, \tag{1}$$

where $\mathbf{D}_j$ and $\mathbf{D}_t$ are a set of age and time dummies. The underlying assumption in the time-effects specification is that changes in $m_{j,t,c}$ over time are due to time-varying factors that affect every age (cohort), and once we control for time effects, we recover the age profiles. Equation (1) is estimated separately for each gender (men and women), marital status (married and single), and skill group. For skills, we divide individuals in two groups; *skilled (s)*, those with at least four years of college education or more, and *unskilled (u)*, with strictly less than college education. The age profiles are given by the estimated β_j values. Then, we also estimate a cohort-effects specification, given by

$$m_{j,t,c} = \beta_j' \mathbf{D}_j + \beta_c' \mathbf{D}_c + \upsilon_{j,t,c}, \tag{2}$$

where $\mathbf{D}_c$ is a set of cohort dummies. In contrast to equation (1), the underlying assumption in the cohort-effects specification is that changes in $m_{j,t,c}$ over time reflect differences between younger and older cohorts.

FIGURE 1.—Variance of log wages, males (left) and females (right).

The assumptions behind these two specifications might be too strong for any specific age profile we construct in this section. In our benchmark analysis, we use the time-effects specification, which accounts better for observed trends in inequality (Heathcote, Storesletten, and Violante (2005)). As a result, for all age profiles we document in this section, we assume that changes over time are captured by time effects that influence all ages at a given time. For some variables, for example, female labor force participation (Goldin (2021)), this assumption might be restrictive, and a cohort perspective might be more appropriate. In Section 7 and in the Supplemental Material, we present our analysis and the profiles with cohort effects.

The key findings that emerge from the benchmark analysis are listed below.

1. For males, as it is well known in the literature, the variance of log hourly wages increases non-trivially along the life cycle; see Figure 1 (left panel). As dummies, estimated β_j values only capture variance of log wages for each age relative to an omitted one. Therefore, we normalize β_{25} to its value in the data (the variance of log wages for 25-year-olds, averaged across years) and rescale all other coefficients accordingly. The increase is more pronounced for skilled than for unskilled men. The increase is of nearly 30 log points for married skilled men between ages 25 and 60, versus a corresponding increase of about 14 log points for married unskilled men. These patterns hold for single men as well, albeit with a smaller increase in variances over the life cycle. These patterns are mirrored when inequality in labor earnings rather than hourly wages is considered.

2. For females, married or single, we *do not observe* a similar increase. This is largely independent of marital status and skill—see Figure 1 (right panel). The increase in dispersion in hourly wages for unskilled (skilled) females is of about five (ten) log points up to age 40, and after that, the level of dispersion is roughly *constant*. This is in stark contrast with the increase in dispersion for males discussed in point 1 above.[4]

3. For both married and single households, the variance of log earnings increases non-trivially along the life cycle, but the level of inequality is *much lower* among married households. At age 25 (45), variance of log earnings is about 0.38 (0.55) for all households, but only 0.30 (0.41) for married households.

4. The gender wage gap, defined as the ratio of average hourly earnings of females relative to males, increases over the life cycle. These changes are sharper for skilled in-

[4]Bayer and Kuhn (2019) documented similar gender differences in life-cycle profiles of earnings inequality in Germany.

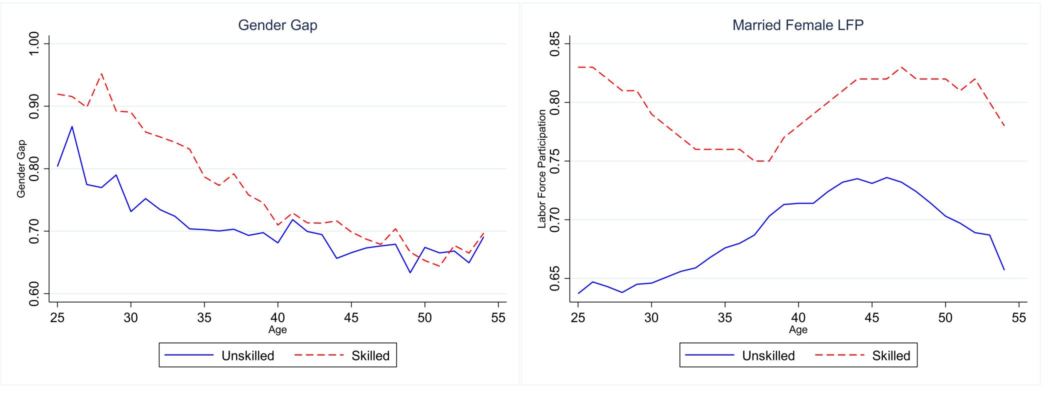

FIGURE 2.—The gender wage gap (left); LFP of married females (right).

dividuals, with a decline in this ratio from about 92% at age 25 to about 70% at age 45. For unskilled individuals, the corresponding change is smaller and of about 11 percentage points. Figure 2 (left panel) displays these patterns.

5. Over the life cycle, the participation rate of married females first declines and then rises up to ages 45–48, and then declines again. These changes are much more pronounced for married *skilled* females. Figure 2 (right panel) displays these patterns.

6. Conditional on work, there is significant variation in hours of work among married females, measured by the variance of log hours at each age. The level is, nevertheless, roughly constant over the life cycle, at around 0.13; see Figure 3 (left panel).

7. The correlation between earnings of husbands and wives is low, around 0.15 at ages 40–50, and slightly ∩-shaped early in the life cycle. Figure 3 (right panel) displays these patterns.

8. The variance of log consumption increases along the life cycle, but *much less than* the increase in the variance of household or individual earnings. The increase peaks at age 55, about 0.12 log points above its level at age 25. This is a well-known fact by now, and was documented in Aguiar and Hurst (2013) and Primiceri and van Rens (2009), among others.

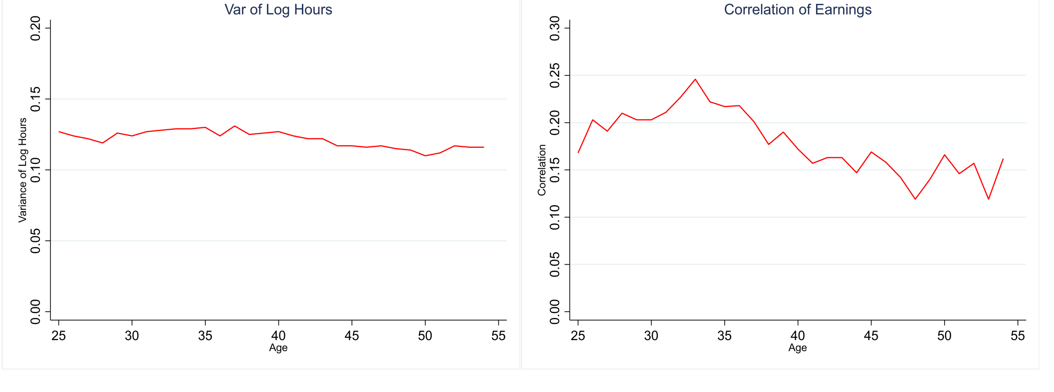

FIGURE 3.—Variance of log hours, married females (left); correlation of spousal earnings (right).

3. THE ECONOMIC ENVIRONMENT

We study a stationary life-cycle economy populated by a continuum of males (m) and a continuum of females (f). Let $j \in \{1, 2, \ldots, J\}$ denote the age of each individual. Each model period is one year, and the first model period corresponds to age 25. Population grows at rate n. The life cycle of agents is split into two parts. Each agent starts life as a worker and, at age J_R, individuals retire and collect pension benefits until they die at age J.

Individuals differ in their marital status. We assume that they are born as either *single* or *married* and their marital status does not change over time. Each individual is also born with a given intrinsic type (education) that defines the rental rate for his/her labor services, and the growth of their labor endowment as they age. Married households are composed of individuals who are of the same age.

Married households and single females also differ in terms of the number of children attached to them. They can be childless or endowed with children. The number of children depends on the educational attainment of the parents and they appear either early or late in the life cycle exogenously. Children affect the resources available to households for several periods, and this is mitigated partially or fully by government policies targeted to children. Children do not provide any utility.

Individuals also differ in terms of permanent shocks received at the start of life, which are correlated among spouses. Furthermore, each period, individuals experience uninsurable productivity shocks, which affect how much they can earn per hour. We assume that these shocks are persistent. We also assume that shocks that husbands and wives receive are correlated. Hence, heterogeneity among households arises due to different factors: their education level, the permanent and life-cycle shocks of their members, and who is married with whom. These forms of ex post and ex ante heterogeneity determine, in conjunction with labor supply and savings decisions, the degree of income, consumption, and wealth inequality in the economy.

Production and Markets. There is an aggregate firm that operates a constant returns to scale technology. The firm rents capital and skilled and unskilled labor services from households at the rates R, w_s, and w_u, respectively. Using K units of capital and L units of the composite labor input, the firm produces

$$F(K, L) = K^{\alpha} L^{1-\alpha}, \quad \text{with } L \equiv \left(\xi L_s^{\rho} + (1-\xi) L_{u,g}^{\rho}\right)^{\frac{1}{\rho}}, \ \rho \in (-\infty, 1),$$

where L_s and $L_{u,g}$ stand for the stock of skilled labor and unskilled labor used in the production of goods, respectively. The elasticity of substitution between labor of different types is constant and given by $\sigma = \frac{1}{1-\rho}$.

We assume that capital depreciates at rate δ_k. Childcare services are provided using unskilled labor services only. Thus, the price of childcare services is the wage rate, w_u. As a result, unskilled labor services available are split between the production of consumption and investment goods, $L_{u,g}$, and childcare services, $L_u - L_{u,g}$. Households save in the form of a risk-free asset that pays the competitive rate of return $r = R - \delta_k$.

Ex Ante Heterogeneity and Demographics. At the start of life, each male is endowed with an exogenous type z that remains constant over his life cycle: $z \in Z = \{u, s\}$. This type of heterogeneity defines whether the agent is *skilled* (s) or *unskilled* (u) that we later map to educational levels in the data. For females, we equivalently have $x \in X = \{u, s\}$.

We assume that each cohort is $1+n$ bigger than the previous one. These demographic patterns are stationary so that age-j agents are a fraction μ_j of the population at any point in time. The weights are normalized to add up to 1, and obey the recursion, $\mu_{j+1} = \mu_j/(1+n)$.

3.1. *Labor Efficiency Units*

We consider a general structure, where individuals differ at the start of the life cycle in their skills, permanent shocks, as well as uninsurable shocks experienced as they age. These shocks are dependent on the skill of individuals (u, s), their gender (m, f), and their marital status (M, S).

Singles. Consider first single males. Their labor endowment (efficiency units) at age j is given by

$$\varpi_m(z, j)\exp\bigl(\nu^S_{m,z} + \eta^S_{m,z,j}\bigr), \quad z \in Z = \{u, s\},$$

where the function $\varpi_m(\cdot,\cdot)$ summarizes the combined effects of skill and age on the labor endowment. ν is a *permanent* shock and η is a *persistent* shock. We assume that the permanent shock is normally distributed: $\nu^S_{m,z} \sim N(0, \sigma^2_{\nu^S_{m,z}})$, $z \in Z$.

We assume that for $j > 1$, the persistent shock is governed by a random walk, given by

$$\eta^S_{m,z,j+1} = \eta^S_{m,z,j} + \varepsilon^S_{m,z,j+1}, \quad z \in Z,$$

with $\varepsilon^S_{m,z,j+1} \sim N(0, \sigma^2_{\varepsilon^S_{m,z}})$ representing innovations over time. We furthermore assume that the initial value of η at the start of the life cycle is zero; that is, $\eta^S_{m,z,1} = 0$, $z \in Z$.

The structure is different for single females, as their efficiency units evolve endogenously, with growth and depreciation rates that depend on intrinsic skills and labor market experience. Intrinsic skills determine their initial human capital: $h_1 = \varpi_f(x, 1)$, $x \in X$. For $j > 1$, we have

$$h' = \mathcal{H}(x, h, l, e) = \exp\bigl[\ln h + \alpha^e_x \chi(l) - \delta_x\bigl(1 - \chi(l)\bigr)\bigr], \quad x \in X = \{u, s\}, \tag{3}$$

where e stands for labor market experience and $\chi(\cdot)$ is an indicator function that is 1 if hours worked are positive and zero otherwise. The parameter α^e_x is the experience-skill growth rate and δ_x stands for the depreciation rate. It follows that for a single female of age j who has human capital h, her realized labor efficiency is given by $h \times \exp(\nu^S_{f,x} + \eta^S_{f,x,j})$. The permanent and the persistent shock obey the same representation as for males, with innovation variances that depend on marital status and skill.

Married Couples. Married individuals draw permanent shocks at the start of their life cycle that are potentially *correlated*. They also draw values for their persistent shocks which are potentially correlated as well. The labor endowments (labor efficiency) of a married male and a married female are given by

$$\varpi(z, j) \times \exp\bigl(\nu^M_{m,z} + \eta^M_{m,z,j}\bigr) \quad \text{and} \quad h \times \exp\bigl(\nu^M_{f,x} + \eta^M_{f,x,j}\bigr), \quad z \in Z, x \in X.$$

The labor efficiency of a married female is correspondingly given by $h \times \exp(\nu^M_{f,x} + \eta^M_{f,x,j})$, where h follows the same law of motion for singles; equation (3).

The initial conditions are such that $\eta^M_{m,z,1} = 0$ and $\eta^M_{f,x,1} = 0$. For $j > 1$, $\eta^M_{m,z,j}$ and $\eta^M_{f,x,j}$ follow a bivariate process, given by

$$\eta^M_{m,z,j+1} = \eta^M_{m,z,j} + \varepsilon^M_{m,z,j+1} \quad \text{and} \quad \eta^M_{f,x,j+1} = \eta^M_{f,x,j} + \varepsilon^M_{f,x,j+1} \quad \text{for } z \in Z, x \in X$$

with

$$\left(\varepsilon^M_{m,z,j+1}, \varepsilon^M_{f,x,j+1}\right) \sim N\left(\begin{matrix} 0 \\ 0 \end{matrix}, \begin{matrix} \sigma^2_{\varepsilon^M_{m,z}} & \sigma_{\varepsilon_f \varepsilon_m} \\ \sigma_{\varepsilon_f \varepsilon_m} & \sigma^2_{\varepsilon^M_{f,x}} \end{matrix}\right), \quad z, x \in Z \times X.$$

The values of permanent shocks for married individuals are draws from a bivariate normal distribution as well. That is,

$$\left(\nu^M_{m,z}, \nu^M_{f,x}\right) \sim N\left(\begin{matrix} 0 \\ 0 \end{matrix}, \begin{matrix} \sigma^2_{\nu^M_{m,z}} & \sigma_{\nu_f \nu_m} \\ \sigma_{\nu_f \nu_m} & \sigma^2_{\nu^M_{f,x}} \end{matrix}\right), \quad z, x \in Z \times X.$$

Note that we assume that while innovations depend on skills, the covariance structure for both permanent and persistent shocks does not. This parsimonious specification allows us to capture key correlations across married spouses, both at the start as well as in along the middle of the life cycle—see Section 5.

Labor Earnings. We now summarize the notion of labor *earnings* resulting from our choices, taking into account skill prices (w_s and w_u), endowments, and labor supply choices—described later. For an age-j single male of type z, earnings are given by

$$\underbrace{w_z}_{\text{wage by skill}} \underbrace{\varpi(z, j) \exp\left(\nu^S_{m,z} + \eta^S_{m,z,j}\right)}_{\text{labor efficiency}} \underbrace{l_m}_{\text{labor supply}}.$$

For a single female of skill $x \in X$ who has human capital h, age j, earnings are given by

$$\underbrace{w_x}_{\text{wage by skill}} \underbrace{h \exp\left(\nu^S_{f,x} + \eta^S_{f,x,j}\right)}_{\text{labor efficiency}} \underbrace{l_f}_{\text{labor supply}}.$$

Finally, for a married couple of skill $z, x \in Z \times X$, of age j, when she has h units of human capital, earnings are given by

$$\underbrace{w_x}_{\text{wage by skill}} \underbrace{h \exp\left(\nu^M_{f,x} + \eta^M_{f,x,j}\right)}_{\text{labor efficiency}} \underbrace{l_f}_{\text{labor supply}} + \underbrace{w_z}_{\text{wage by skill}} \underbrace{\varpi(z, j) \exp\left(\nu^M_{m,z} + \eta^M_{m,z,j}\right)}_{\text{labor efficiency}} \underbrace{l_m}_{\text{labor supply}}.$$

3.2. *Children and Childcare Costs*

Children are assigned exogenously to married couples and single females at the start of life, depending on the education of parents. Each married couple and single female can be of three types: *without* any children, *early* child bearers, *late* child bearers. We denote this dimension of heterogeneity by $b = \{0, 1, 2\}$.

If $b \neq 0$, children appear deterministically at parents' age $\bar{j}(x, z, b)$ for married households and $\bar{j}(x, b)$ for single females. Married households have $k(x, z)$ children, while single females have $k(x)$ children. For married households, half of the children appear at age $\bar{j}(x, z, b)$ and the other half at age $\bar{j}(x, z, b) + 2$; that is, children are spaced by two years. It is equivalent for single households: half of the children appear at age $\bar{j}(x, b)$ and the other half at age $\bar{j}(x, b) + 2$. Each child stays with their parents for N model periods.

We assume that if a female with children works, married or single, then the household has to pay for childcare costs. Childcare costs depend on the age of the child, t, and are priced at rate w_u. We assume that children in single female households require $d(x, t)$ units of childcare services per child, $t = 1, 2, \dots, N$. Married households require $d(x, z, t)$ units of childcare services per child. Since competitive price of childcare services is the unskilled wage rate w_u, the cost of childcare services per child equals $w_u d(x, t)$ for single females and $w_u d(x, z, t)$ for married households.

3.3. *Preferences*

The momentary utility function for singles is given by

$$U^S(c, l) = \log(c) - B_i(l)^{1+\frac{1}{\gamma}}, \quad i = m, f,$$

where c is consumption, l is time devoted to market work, and γ is the intertemporal elasticity of labor supply (Frisch elasticity). The parameter B_i captures potential gender-driven differences in the disutility of work.

Married households maximize the sum of their members' utilities. We assume that when the female member of a married household works, the household incurs a utility cost q. We assume that at the start of their lives, married households draw a $q \in Q$, where $Q \subset R_{++}$ is a finite set. These values of q represent the utility costs of joint market work for married couples. For a given household, the initial draw of utility cost depends on the type (education) of the husband. Let $\zeta(q|z)$ denote the probability that the cost of joint work is q, with $\sum_{q\in Q} \zeta(q|z) = 1$. We assume that for married households with children at home, the utility cost q is multiplied by a factor that depends on the age of the youngest child at home, $t_{\min}$, and the mother's skill level, $\vartheta_x(t_{\min})$, $x \in X$. This specification captures the idea that joint work becomes more costly with arrival of children, beyond childcare costs, and that this additional cost changes as children grow older.

Formally, if $b \in \{1, 2\}$ and the household age is such that $\bar{j}(x, z, b) \le j \le \bar{j}(x, z, b) + N + 2$, that is, children are at home (recall that the first child arrives at $\bar{j}(\cdot)$ and the second one leaves at $\bar{j}(\cdot) + N + 2$), then the period utility of a married household is given by

$$U^M(c, l_f, l_m; \theta, q, j) = 2\log(c) - B_m l_m^{1+\frac{1}{\gamma}} - \theta B_f l_f^{1+\frac{1}{\gamma}} - \chi\{l_f\} q\big(1 + \vartheta_x(t_{\min})\big), \qquad (4)$$

where $\chi\{\cdot\}$ denotes the indicator function.[5] For households without any children at home, $\vartheta_x(t_{\min}) = 0$.

Note that consumption is a public good within the household. The variable θ captures heterogeneity in the disutility of work across married females. We assume that θ is realized at the start of life, and takes two values with equal probability: $\theta \in \{\theta_L, \theta_H\}$. Note also that the parameter $\gamma > 0$, the intertemporal elasticity of labor supply, is common for all individuals: males or females, married or single. It is also important to note that following

[5]Note that if x, z, and j are known, the age of the youngest child can be readily calculated.

the tradition in macroeconomics literature, we restrict the preferences to be consistent with a balanced-growth path. As in, for example, Attanasio, Banks, Meghir, and Weber (1999, 2005), we could allow the marginal utility of consumption to be affected by the female labor force participation decision. In the current specification, the female labor force participation and demographics (the number of children) affect the level of utility through the cost of joint work.

3.4. *Taxes and Transfers*

There is a government that taxes labor and capital income, and uses tax collections to pay for government consumption, tax credits, and transfers to individuals. It also runs a pay-as-you-go social security system, so it collects payroll taxes and pays retirement benefits.

Transfers. Households in the model have access to transfers that depend on gender, marital status, and household income. Income for tax and transfer purposes is labor plus asset income. For a household with income level I, number of children k, and child-care expenses D, the transfers are represented by functions $TR_f^S(I, k, D)$, $TR_m^S(I)$, and $TR^M(I, k, D)$, for a single-female, single-male, and married-couple household, respectively. This generic formulation of transfers allows us to capture a host of transfers and tax credit programs in the United States. We describe below how these functions are parameterized in light of data.

Taxation and Social Security. The total income tax liabilities of married and single households, before any tax credits, are affected by the presence of children in the household, and are represented by tax functions $T^M(I, k)$ and $T^S(I, k)$, respectively, where k stands for the number of children at the household. These functions are continuous in I, increasing, and convex. This representation captures the effective variation in tax liabilities associated to income, marital status, and the presence of children in households.

There is a (flat) payroll tax that taxes individual labor incomes, represented by τ_p, to fund social security transfers. Moreover, each household pays an additional flat capital income tax for the returns from his/her asset holdings, denoted by τ_k. Retired households have access to social security benefits. The social security benefits depend on agents' education types, that is, initially more productive agents receive larger social security benefits. This allows us to capture in a parsimonious way the positive relation between lifetime earnings and social security transfers, as well as the intra-cohort redistribution built into the system. Let $p_f^S(x)$, $p_m^S(z)$, and $p^M(x, z)$ indicate the level of social security benefits for a single female of type x, a single male of type z, and a married retired household of type (x, z), respectively. The social security system has to balance its budget every period.

3.5. *Decision Problem*

We now present the decision problem for different types of agents in the recursive language. We provide a formal definition of a stationary equilibrium in Guner, Kaygusuz, and Ventura (2023b). We focus on single females and married couples, since the problem of single males is rather standard. For ease of notation, the dependence of shocks on type, gender, and marital status is suppressed whenever possible. For single females, the individual state is $(a, h, e, x, \nu_{f,x}^S, \eta_{f,x}^S, b, j)$, where a stands for asset holdings. For married couples, the state is given by $(a, h, e, x, z, \theta, \nu_{m,z}^M, \nu_{f,x}^M, \eta_{m,z}^M, \eta_{f,x}^M, q, b, j)$.

Note that the dependency of transfers and taxes on the presence of children in the household is summarized by age of parents (j) and childbearing status (b), in conjunction with x for single females and the pair (x, z) for married couples. The same reasoning applies for childcare costs, or the utility costs of joint participation for married couples when children are present. That is, if we know the intrinsic type of a single female or a married household, the age of parents (j), and fertility type (b), we know the age of each child and the childcare costs. Given parents' types, the half of children appear at parents' age $\bar{j}(\cdot)$ and the other half at $\bar{j}(\cdot)+2$. Then, when their parents are of age j, young and old children at home have ages $j-\bar{j}(\cdot)+1$ and $j-\bar{j}(\cdot)+3$.

For expositional purposes, we collapse the permanent/exogenous characteristics in the household problems in a single vector of state variables. For single females, let $\mathcal{S}_f^S \equiv (x, \nu_{f,x}^S, b)$ be the vector of variables that do not change along the life cycle for single females and single males, respectively. For married households, let $\mathcal{S}^M \equiv (x, z, \theta, \boldsymbol{\nu}, q, b)$ be the vector of such states for married households, with $\boldsymbol{\nu} \equiv (\nu_{f,x}^M, \nu_{m,z}^M)$. In similar fashion, for the case of married couples, we summarize the pair of persistent shocks by $\boldsymbol{\eta} \equiv (\eta_{f,x}^M, \eta_{m,z}^M)$. Likewise, for expositional purposes, we denote by $\mathcal{E}_f^S(x, h, \eta_{f,x}^S, \nu, l_f)$ and $\mathcal{E}^M(x, z, h, \boldsymbol{\eta}, \boldsymbol{\nu}, l_m, l_f, j)$ the labor earnings of single females and married couples, respectively, as defined in Section 3.1.

The Problem of a Single Female Household. Given her current state, $(a, h, e, \mathcal{S}_f^S, \eta, j)$, the problem of a single female is

$$V_f^S\big(a, h, e, \mathcal{S}_f^S, \eta, j\big) = \max_{a', l}\big\{U^S(c, l) + \beta \mathbf{E}_{\eta'|\eta} V_f^S\big(a', h', e', \mathcal{S}_f^S, \eta', j+1\big)\big\},$$

subject to (i) *With kids*: if $b = \{1, 2\}$, $j \in \{\bar{j}(x, b), \bar{j}(x, b)+1, \dots, \bar{j}(x, b)+N+2\}$,

$$c + a' = \begin{cases} a\big(1 + r(1-\tau_k)\big) + \mathcal{E}_f^S(x, h, \eta, \nu, l)(1-\tau_p) \\ \quad + TR_f^S(I, \mathcal{K}, D) - T^S(I, \mathcal{K}) - w_u D \chi(l), \end{cases}$$

where $I = \mathcal{E}_f^S(x, h, \eta, \nu, l) + ra$. $\mathcal{K}$ is the number of children present in the household, either old, born at $\bar{j}(x, b)$, or young, born at $\bar{j}(x, b)+2$. It is given by

$$\mathcal{K} = \frac{k(x, b)}{2}\big[\underbrace{\chi\big(\bar{j}(x, b) \le j \le \bar{j}(x, b)+N\big)}_{\text{old children}} + \underbrace{\chi\big(\bar{j}(x, b)+2 \le j \le \bar{j}(x, b)+2+N\big)}_{\text{young children}}\big].$$

Meanwhile, D stands for the childcare expenses incurred:

$$\begin{aligned} D = {} & \frac{k(x, b)}{2} d\big(x, j - \bar{j}(x, b) + 1\big)\chi\big(\bar{j}(x, b) \le j \le N\big) \\ & + \frac{k(x, b)}{2} d\big(x, j - j(x, b) + 3\big)\chi\big(\bar{j}(x, b) + 2 \le j \le \bar{j}(x, b) + 2 + N\big). \end{aligned}$$

(ii) *Without kids but not retired*: if $b = 0$, or $b = \{1, 2\}$ and $j \notin \{\bar{j}(x, b), \dots, N+2\}$, then there are no children at home and

$$\begin{aligned} c + a' = {} & a\big(1 + r(1-\tau_k)\big) + \mathcal{E}_f^S(x, h, \eta, \nu, l)(1-\tau_p) \\ & + TR_f^S(I, 0, 0) - T^S(I, 0). \end{aligned}$$

(iii) *Retired*: if $j \geq J_R$, then there are no children and

$$c + a' = a\big(1 + r(1 - \tau_k)\big) + p_f^S(x) - T^S(ra, 0) + TR_f^S(ra, 0, 0).$$

In addition,

$$\begin{aligned} h' &= \mathcal{H}(x, h, l_f, e) = \exp\big[\ln h + \alpha_e^x \chi(l_f) - \delta^x\big(1 - \chi(l_f)\big)\big], \\ e' &= e + \chi(l) \text{ and } l \geq 0, a' \geq 0 \text{ (with strict equality if } j = J + 1). \end{aligned} \tag{5}$$

The Problem of Married Households. Like singles, married couples decide how much to consume, how much to save, and how much to work. They also decide whether the female member of the household should work, taking into account the evolution of her skills, experience, and childcare costs. Note that in the formulation below, we make the current utility of married households to depend on (x, z, b, j), as these variables fully determine the age of children present in the household that may affect the disutility of joint market work, $q(1 + \vartheta_x(t_{\min}))$ term above. Formally, the problem is given by

$$\begin{aligned} V^M\big(a, h, e, \mathcal{S}^M, \boldsymbol{\eta}, j\big) = \max_{a', l_f, l_m} \big\{&U^M(c, l_f, l_m, q, x, z, b, j) \\ &+ \beta \mathbf{E}_{\boldsymbol{\eta}'|\boldsymbol{\eta}} V^M\big(a', h', e', \mathcal{S}^M, \boldsymbol{\eta}', j + 1\big)\big\}, \end{aligned}$$

subject to

(i) *With kids*: if $b = \{1, 2\}$, $j \in \{\bar{j}(x, b), \dots, N + 2\}$, then

$$c + a' = \begin{Bmatrix} a\big(1 + r(1 - \tau_k)\big) + \mathcal{E}^M(x, z, h, \boldsymbol{\eta}, \boldsymbol{\nu}, l_m, l_f, j)(1 - \tau_p) \\ -T^M(I, \mathcal{K}) + TR^M(I, \mathcal{K}, D) - w_u D \chi(l) \end{Bmatrix},$$

where $I = \mathcal{E}^M(x, z, h, \boldsymbol{\eta}, \boldsymbol{\nu}, l_m, l_f, j) + ra$, and $l_m \geq 0$, $l_f \geq 0$, and $a' \geq 0$.

In this formulation, $\mathbf{E}_{\boldsymbol{\eta}'|\boldsymbol{\eta}}$ now represents the joint expectation over the shocks that husbands and wives face. The number of children present, $\mathcal{K}$, and childcare expenses, D, are formulated as they were done for a single female. The household problem also takes into account the accumulation of human capital for the wife, given by $h' = \mathcal{H}(x, h, l_f, e)$ and $e' = e + \chi(l)$. The budget constraints when the household is not retired but without any children and when the household is retired, cases (ii) and (iii), are defined accordingly.

3.6. *Sources of Inequality in the Model*

What are the determinants of inequality at a point in time and over the life cycle across individuals and households in the model? This question is of central importance in assessing the effects of transfer policies.

First, individuals differ in their intrinsic skills and experience permanent and persistent shocks. Permanent and persistent shocks are common in life-cycle models with heterogeneous individuals. Different from most of the work in the area, differences in skill type at birth determine (i) potentially different growth rates in labor productivity between skilled and unskilled individuals, and (ii) between-group differences as individuals face different rental rates for labor services depending on their skill type. Point (i) implies that our model encompasses a mixture of traditional parameterization of heterogeneity (usually referred to as Representative Income Processes or RIP), with a human capital view of differences of individuals as they age, as emphasized in Guvenen (2009) and Huggett, Ventura, and Yaron (2011), among others (Heterogeneous Income Processes or HIP).

The second layer of heterogeneity determining inequality concerns marital status. At birth, some individuals are single, some are married, and married ones are assigned to spouses according to their skill type. Besides, within a given skill pair, permanent and persistent shocks are potentially correlated between spouses. Overall, as in Greenwood, Guner, Kocharkov, and Santos (2016) and others, marriage can amplify existing differences between individuals and contribute to propagating shocks over the life cycle.

Finally, differences in individuals by gender, coupled with children's presence, help define the level of gender premia in wages at birth and its evolution over the life cycle. As children appear and women leave the workforce, skill depreciation kicks in, and thus, the gender gap in wage rates grows over time. As children require fewer resources as they age, some women return to work, accumulate skills again, and the gender-wage gap moderates its growth. As we describe below in our analysis of the benchmark economy, women's behavior regarding participation over time, in conjunction with uninsurable shocks, determines gender differences in the life-cycle profile of earnings inequality.

The reforms of the welfare state that we consider, the Universal Basic Income (UBI) and Negative Income Tax (NIT), have simple structures. However, they replace the existing welfare state, which is not simple at all. As a result, identifying the winners and losers requires rich ex ante heterogeneity, since different socioeconomic groups receive quite different transfers in the benchmark economy.

3.7. *Modeling Choices*

Several model elements are taken as exogenous in the current analysis. This allows us to simulate a model with extensive heterogeneity in educational attainment, marital status and sorting, and the number and timing of children. As we emphasize in the paper, such heterogeneity is crucial to understanding the welfare consequences of reforms to the welfare state. We model these features as exogenous, which allows us to focus on a subset of critical endogenous decisions: household labor supply, female human capital accumulation, and savings. There is an extensive empirical literature on the incentive effects of the welfare state on marriage and fertility, with a particular focus on the impact of the 1996 welfare reform. The findings from this literature suggest that the incentive effects of the welfare state on marriage, fertility, and single motherhood are modest; see Bitler, Gelbach, Hoynes, and Zavodny (2004), Kearney (2004), and Moffitt, Phelan, and Winkler (2020). Hence, we expect that the welfare state's direct impact on marriage and fertility incentives is likely to be small. On the other hand, it is well established that children that grow with single mothers receive relatively much less investment than those with two parents. As a result, even small adverse effects on marriage incentives coupled with adverse effects on children can accumulate across generations, impacting intergenerational mobility. This was emphasized by Aiyagari, Greenwood, and Guner (2000), among others.

We also have a unitary model of household decisions. With a non-unitary model like in Voena (2015), transfers can affect within-household allocations even if they do not change marriage and divorce decisions. In particular, a non-unitary model could allow us to study how gender-specific taxes and transfers can alter household allocations and within-household distribution of welfare.

The model also abstracts from health shocks and government-provided health insurance programs prior to retirement, such as Medicaid, Social Security Disability Insurance (SSDI), and Children's Health Insurance Program (CHIP). Overall, since 2015, expenditures on these programs have averaged about 3.9% of GDP, with about 3.1% on

Medicaid, 0.7% on SSDI, and about 0.1% on CHIP.[6] Recent papers introduced exogenous health shocks into heterogeneous-agent macro models—see Hosseini, Kopecky, and Zhao (2021) for a recent example. In these models, health shocks affect earnings (mainly through labor supply) and act like income shocks, and as a result, government-provided health insurance is valued by individuals. Hosseini, Kopecky, and Zhao (2021), for example, found that health shocks account for about 30% of lifetime earnings inequality, and government-provided disability insurance is valued by consumers. Since health shocks are exogenous in these models, we see our paper and this literature as complementary, each focusing in detail on different aspects of the welfare system.

4. PARAMETER VALUES

This section describes how we select parameter values to compute a stationary equilibrium. We relegate details to the Supplemental Material (Tables SA-XI and SA-XII summarize our parameter choices). The model period is one year. Agents start their life at age 25, potentially work for forty years, retire at age 65 ($j = J_R$), and then live until age 80 ($j = J$). The population grows at the annual rate of 1.1%. *Skilled* individuals are those with at least a four-year college degree. The marital structure (who is single, who is married, and who is married with whom), childbearing status, and the number of children for different types of households are taken directly from the data.

Endowments. For males, following the procedure described in Section 2, we construct age profiles of mean hourly wages for each skill group using data from 1980–2019 CPS March Supplement, and set $\varpi_m(z, j)$, $z = u, s$, to these profiles (Figure SA7 in the Supplemental Material). For females, we use age-25 wage levels to calibrate their initial human capital levels, $h_1 = \varpi_f(x, 1)$. After age 25, female skills evolve according to equation (3).

We select the parameter α_x^e so that if a type-x female works for one more period, her wage grows exactly at the same rate as a male of the same type with the same experience level (e). Hence, if a female works in every period, her labor market productivity evolves exactly like a male, except for the observed age-25 wage gender gap. Figure SA7 in the Supplemental Material shows the calibrated values for the growth factors. For depreciation rates, we select each one so that the model is consistent with the evolution of the wage gender gap for the first decade of the life cycle (ages 25–35). The resulting values are $\delta_u = 0.025$, and a non-trivially higher value for skilled females, $\delta_s = 0.059$. These values are required to reproduce the faster increase in the wage gender gap with age for skilled females documented in Section 2.[7]

Productivity Shocks. There are in total *eighteen* parameters that determine the productivity shocks: eight variances for permanent shocks (by skill, gender, and marital status), eight innovation variances for persistent shocks (again by skill, gender, and marital status), plus two covariances (for permanent shocks and innovations to persistent shocks). Table SA-XI in the Supplemental Material presents these parameters. For permanent shocks (ν), we match the observed variances of log wages at age 25 by skill, gender, and marital status. To pin down the value of the covariance term for married individuals, $\sigma_{\nu^f \nu^m}$,

[6]For expenditures on Medicaid and CHIP, see Centers for Medicare and Medicaid Services (https://www.cms.gov/). For expenditures on disability, see Annual Statistical Supplements to the Social Security Bulletin (https://www.ssa.gov/policy/docs/statcomps/supplement/).

[7]Blundell, Costa Dias, Meghir, and Shaw (2016) found similar results for the UK.

we target the correlation in log wages among all spouses at age 25. For the variances of innovations to persistent shocks (ε), we target the observed variances of log wages towards the end of the life cycle (age 54) for each group. For the covariance of innovation in persistent shocks across spouses, $\sigma_{\varepsilon^f \varepsilon^m}$, we target the correlation of wages between husbands and wives by middle age (ages 40–45). Overall, the variances of innovations for persistent shocks for men are substantially larger than for females, while the corresponding variances for skilled individuals, male or female, are larger than for unskilled ones. Overall, not surprisingly, the innovation variances are smaller than in related estimates, for example, Heathcote, Storesletten, and Violante (2010) and Huggett et al. (2011). This reflects the division of individuals between skilled—who experience faster growth in labor efficiency with experience—and unskilled ones, as well as the distinction of individuals by gender and marital status.

Income Taxation. To compute the tax functions, that is, $T^S(I, k)$ and $T^M(I, k)$, we adopt a parametric form for the average tax rate:

$$\tau(I) = 1 - \lambda I^{-\tau},$$

where I (income) is measured in multiples of mean household income and $\tau(I)$ is the average tax rate. The parameter τ determines the progressivity of the tax scheme and λ determines its level. The parameters τ and λ depend on marital status and the number of children, and are estimated from IRS micro data on tax returns. Since the EITC, CTC, and CDCTC are explicitly modeled in the benchmark economy, the tax functions are estimated using tax liabilities before these credits are applied.

Transfers. In the Supplemental Material and Guner, Kaygusuz, and Ventura (2023b), we provide a description of the various means-tested programs in the United States, focusing on who qualifies for them and how household's marital status and number of children affect access.[8] We divide these programs into three groups: (i) the Earned Income Tax Credit (EITC); (ii) child-related transfers, which encompass the Child Tax Credit (CTC), the Child and Dependent Care Tax Credit (CDCTC), and childcare subsidies; and (iii) the amalgam of programs that provide cash or in-kind transfers that are routinely identified as the "welfare system," such as the Temporary Assistance to Needy Families (TANF) and the Supplemental Nutrition Assistance Program (SNAP). We calculate that expenditures in all these programs at all levels amounted to about 2.3% of GDP in 2019.[9]

In computing the transfer functions $TR^S_f(I, k, D)$, $TR^S_m(I)$, and $TR^M(I, k, D)$, the main object of this paper, we model tax credits exactly as they appear in the tax code. Following the discussion in Guner, Kaygusuz, and Ventura (2020), the government covers 75% of the childcare costs for households whose income is below a threshold. We chose the threshold so that the poorest 5% of children receive the subsidy.

The final component is the means-tested transfers. Following Guner, Kaygusuz, and Ventura (2020), we use data from the Survey of Income and Program Participation (SIPP)

[8]More extended discussions can be found, among others, in Moffitt (2003) and Guner, Kaygusuz, and Ventura (2020).

[9]In 2019, the U.S. federal government spent 361 billion dollars for non-medical means-tested transfer programs for the working-age population. This is about 8% of the total federal budget and corresponds to about 1.7% of the U.S. GDP. The total spending, federal and state level, amounted to about 2.3% of the GDP and is expected to grow as we write. Total spending at all levels is calculated based on information from Rector and Menon (2018).

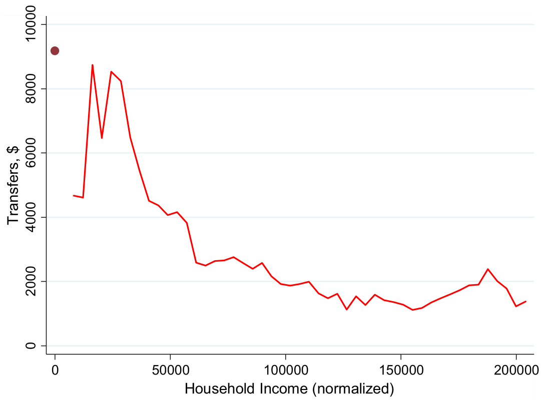

FIGURE 4.—Total transfers in the benchmark economy.

to estimate an effective transfer schedule that relates transfers received by different household types to their income. The welfare payments include all the main means-tested programs as described in the Supplemental Material. We assume that these transfers as a function of income take following form:

$$W(I) = \begin{cases} \omega_0 & \text{if } I = 0, \\ \max\{0, \omega_1 - \omega_2 I\} & \text{if } I > 0, \end{cases}$$

where ω_0 is the transfers for a household with zero income and ω_2 is the benefits reduction rate. Our estimates show that a single female with two children receives about 12% of mean household income in the economy in terms of welfare transfers (about \$12,000 in 2019). Transfers decline gradually with income and vanish at around 1.1 times mean income for a single female with two children (about \$108,000 in 2019). A single female with two children and half of mean household income (about \$44,000 in 2019) receives about \$5800 per year. A married couple with two children who has zero income gets about \$8800. Transfers decline to zero, as they do for a single mother, at around 1.1 times the mean income.

Figure 4 shows how the total transfers (the sum of these three components) vary by household income in the benchmark economy. Households without any income receive transfers in excess of \$8000 per year. The transfers decline sharply for household with positive but very low income. After that, transfers bounce back to around \$8000 and decline smoothly with household income and amount to about \$1000 for households with 1.5 times the mean household income in the economy.

Childcare Costs. To determine the requirement of efficiency units for childcare, $d^M(x, z, t)$ and $d^S(x, t)$, we use data on total spending (as a fraction of household income) on childcare and the relation between children's age and childcare spending (as shown in Figure SA4 in the Supplemental Material). In particular, we use data from the Survey of Income and Program Participation (SIPP), and estimate a relationship between spending in childcare per child and the average age of children, conditional on the mother's skill and marital status. Given the price of unskilled labor services, we recover the efficiency units required at each age in stationary equilibrium.

Remaining Parameters. We select the remaining parameters to match jointly several targets. (i) We set the Frisch elasticity parameter γ to 0.2, and given γ, select the parameters B_m and B_f to match average market hours per worker by gender. The disutility of work shocks is specified as $\theta_L = \exp(-\Delta)$ and $\theta_H = \exp(+\Delta)$, and $\triangle$ is set so as to reproduce the observed variance of log hours of married females at age 40. We choose the discount factor to match capital-to-output ratio (2.9) (ii) We parameterize the distribution of the disutility of joint market work, $\zeta(q|z)$, as a Gamma distribution and infer its parameters to generate the observed female force participation by married females conditional on the husbands' types. Given $\zeta(q|z)$, we determine the loading factors $\vartheta_x(t_{\min})$ so that the model is consistent with the participation rate of mothers by the age of their youngest child present at home (shown in Figure SA4 in the Supplemental Material). (iii) We set the capital share to $\alpha = 0.343$ and the depreciation rate of capital to $\delta^k = 0.055$. To select the parameter governing the elasticity of substitution, ρ, we use standard estimates of this elasticity that suggest a value of 1.5—see, for example, Katz and Murphy (1992). This dictates $\rho = 1/3$. To calibrate the share parameter ξ, we force the model to reproduce the aggregate *skill premium* in the data, defined as per-worker earnings of workers in the skilled category to per-worker earnings of workers in the unskilled category. For this statistic, we target a value of 1.8. (iv) Finally, we pick the additional proportional tax, τ_k, on capital so that the model matches corporate tax collections from data, and select the social security benefits, b, for a given tax rate from the U.S. data, to balance the social security budget.

5. THE BENCHMARK ECONOMY

In Table I, we show summary statistics on how the model performs regarding targeted and non-targeted moments. Total transfers in the model are about 2.3% of the GDP, which (endogenously) matches the data counterpart. The model reproduces the growth in dispersion in hourly wages for married individuals by skill, the correlation of wages of married couples at the start and the end of the life cycle, and married females' participation rates. Differently from other papers in the literature, the model is in line with skill premia. Among other factors, this is driven by the fact that rental rates for labor services differ by skill as skilled and unskilled efficiency units are not perfect substitutes in production.

Importantly, the model is in line with the (non-targeted) initial level and growth in household consumption dispersion over the life cycle (Figure A8). The growth in consumption inequality is lower than the growth in earnings dispersion for males or for households, as we noted in Section 2. One reason for this finding is that several factors contributing to dispersion in earnings with age are anticipated as of the start of the life cycle (an aspect emphasized by Huggett et al. (2011)). Furthermore, transfers contribute to total household income at the bottom of the income distribution and the total household income matters for consumption inequality. Finally, labor supply adjustments along the intensive and extensive margin help households smooth idiosyncratic wage shocks.

The bottom panel of Table I shows earnings inequality measures in the model and the data for households with heads between ages 25 and 65. The model captures the 90–10 and 90–50 ratios very well, and is able to produce earnings shares of the bottom 10%, 20%, and 40% of households, which is critical for the analysis at hand.[10] Not surprisingly,

[10]Inequality measures in the data are from 2010 CPS, with sample restrictions as detailed in the Supplemental Material. As pointed out by Kuhn and Rios-Rull (2016), there is also significant inequality among single

FIGURE 5.—Variance of log wages, model versus data, males (left), females (right).

taxes and transfers reduce income inequality nontrivially; for ages 25–64, the 90–10 ratio is 6.8 for after-transfer household income and 6.2 for after-tax-and-transfer household income. For before-tax-andtransfer household income, the ratio is 7.9.

In sum, the benchmark economy is able to generate observed inequality in the data, it matches the total spending in transfers as a fraction of GDP when various transfer programs are modeled as closely as possible to how they operate, and it delivers the level and the growth in consumption inequality observed in the data. Altogether, these features suggest that the model economy is a suitable framework for studying the welfare state.

Life-Cycle Statistics. Our model environment is consistent with a host of observations over the life cycle. We start by noting that our economy generates the observed growth in dispersion in hourly wages by skill, gender, and marital status. Figure 5 illustrates this. We now concentrate on three interconnected life-cycle statistics. First, we note that our economy generates the life-cycle pattern of the wage gender gap, as Figure 6 (left panel) demonstrates. The model, parameterized to generate the decline in the gender gap by skill in the early ages of the life cycle, captures quite well the slow opening of the gap for unskilled workers over the entire life cycle. The model generates the gradual opening of the gap for skilled workers but leaves a portion unaccounted for towards the end of the working life cycle. At age 50, skilled females earn 66% on average relative to men in the data, while the model implies a gender gap of 76%.

Figure 6 (right panel) shows the performance of the model regarding participation rates of married females as they age. The reader should recall that the economy is parameterized to reproduce the aggregate levels of participation rates by household type, and their levels as of age 40. The endogenous forces inside the model—costly children and utility costs of joint participation that vary with the age of children—lead to the horizontal S-like pattern of participation rates of married females in the data, as the figure demonstrates. The model environment also captures well the initial rise and slow decline of unskilled married females. Overall, this leads the model economy to reproduce well the aggregate pattern of participation rates as individuals age.

---

and married households as well. The model delivers substantial within-group heterogeneity for married and single households. The 90–10 earnings ratios for married and single households are 5.5. and 6 in the model. The 90–50 ratios are 2.3 and 2.4, respectively.

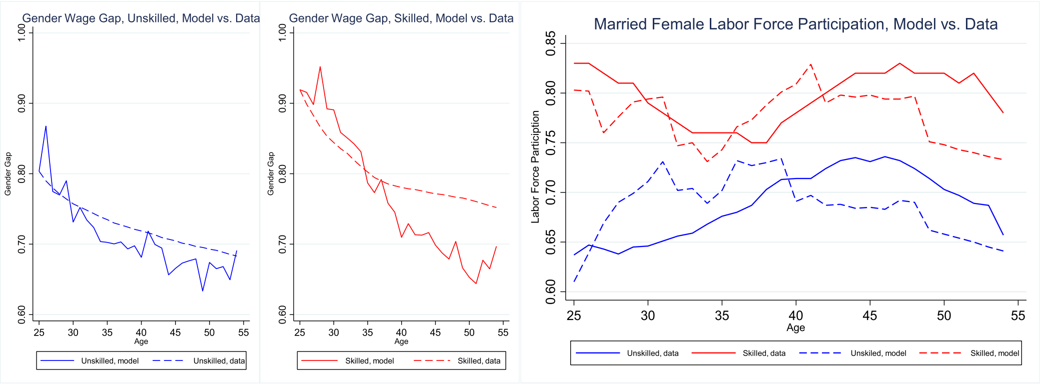

FIGURE 6.—Gender wage gap (left); LFP of married females (right), model versus data.

Overall, as a result of all the forces of our economy operating in tandem, our model implies an age pattern of dispersion in earnings for married females that is broadly consistent with observations. Recall from Section 2 that dispersion in wages of married females first rises, and unlike the case of men, it flattens out as of age 35. As Figure 5 shows, our model generates the same patterns. Why? Early in the life-cycle, skilled females increase their skills faster as a group relative to their unskilled counterparts. This, in conjunction with life-cycle shocks, leads to the overall increase in earnings inequality. In the meantime, some women gradually return to work—given the gradual reduction in childcare and utility costs of joint participation as children age—and start increasing their skills by acquiring experience. Since their skills are lower but accumulate faster, inequality first grows but subsequently starts leveling off. Eventually, all differential rates in skill formation become less and less important as individuals age, and females become more homogeneous. The net result is a flat profile of earnings dispersion after middle age, as the figure shows.

Children and Childcare Costs. What is the quantitative importance of children and childcare costs? To answer this question, we set all childcare costs to zero, while keeping all other parameters constant. We find that childcare costs matter critically in determining the levels of participation rates, and how inequality in wages and earnings evolve over the life cycle for married females. When childcare costs are set to zero, the participation rate of unskilled married females is 74.7%, while for skilled, it is 81.4%. The values in the benchmark model are 68.7% and 77.7%, respectively. The model cannot generate the observed sharp decline in labor force participation early in the life cycle, demonstrating it is associated with childcare requirements. Furthermore, without children, the variance of log wages grows linearly along the life cycle for women, exactly as it does for men (see Figure SA9 in the Supplemental Material).

5.1. *How Valuable Is the Welfare State?*

How much do households value the current transfer scheme? What would be the effects of abolishing the welfare state? To answer these questions, we proceed by fully eliminating all transfers comprising the welfare state. We balance the budget by adjusting the ‘level’ parameter of the tax function (λ) in a proportional and symmetric way for all households. Further, as in all subsequent experiments that we conduct, we assume that the rate of

TABLE I

MODEL AND DATA.

Aggregates	*Data*	*Model*
Capital Output Ratio	2.9	2.9
Total Transfers (% of GDP)	2.3	2.3
Skill Premium	1.8	1.8
LFP of Married Females (%), 25–54		
Unskilled	68.3	68.7
Skilled	77.6	77.7
Total	71.5	72.3
Life-Cycle Inequality		
Variance log wages (Married Males, age 54, S)	0.45	0.45
Variance log wages (Married Males, age 54, U)	0.34	0.34
Variance log wages (Married Females, age 54, S)	0.35	0.35
Variance log wages (Married Females, age 54, U)	0.26	0.26
Variance log hours (Married Females, age 40)	0.13	0.13
Correlation Between Wages of Spouses (age 25)	0.31	0.31
Correlation Between Wages of Spouses (age 40)	0.34	0.33
Variance log consumption (Age 54 vs 25)	0.12	0.12
Earnings Inequality (25–64)		
90–10 ratio	7.8	7.1
90–50 ratio	2.6	2.5
Share, bottom 10%	1.8	2.1
Share, bottom 20%	4.5	5.5
Share, bottom 40%	13.2	15.8

Note: Entries summarize the performance of the benchmark model in terms of empirical targets and key aspects of data. The data for aggregate inequality statistics take into account the same data restrictions used in the empirical analysis in Section 2.

return on capital does not change across steady states. In Guner, Kaygusuz, and Ventura (2023b), we also discuss the elimination of different transfers one at a time.

Table II presents the main findings. Hours worked increase across the board, and these increases are concentrated among the unskilled. The participation rate of married females increases by 6.3% for unskilled women and by 2.0% for skilled ones. All this contributes to a total increase in labor hours of about 3.0% and an increase in aggregate output of 1.7%. The average tax rate at mean income falls substantially, by more than four percentage points across the board: from about 9.2% to 4.8% for a married household with two children, and from about 7.7% to about 3.3% for a single female with two children.[11] When transfers are eliminated, labor supply increases for low- and middle-income households. The increase in labor supply is partly due to income effects and partly due to the incentives to increase labor supply for insurance motive. These changes occur despite the removal of programs that provide incentives for labor supply (e.g., childcare subsidies via the CCDF) or include provisions that subsidize work (e.g., EITC). For all households, the work incentives also increase because of lower taxes.

Table II also shows large negative effects on aggregate welfare, with a compensating variation of about −3.2% for all newborn households. Benchmark transfers are substantial and concentrated at the bottom of the skill distribution. Hence, their elimination leads to significant welfare losses in a utilitarian sense. Single females bear the brunt of the transfer elimination. A newborn, single unskilled female experiences a loss of 5.2%

[11]To balance the budget in this exercise, we multiply λ values in Table SA-VI in the Supplemental Material by 1.048 (recall that $1 - \lambda$ is the tax rate at the mean income).

TABLE II

ELIMINATING ALL TRANSFERS.

Aggregates		Welfare	
(% changes relative to benchmark)		(Newborns, %)	
Output	1.7	*Single F*	
Aggregate Hours	3.0	Unskilled	−5.2
Hours per worker (All Females)	3.2	Skilled	−1.2
Hours per worker (All Males)	1.9	*Married*	
		Unskilled, Unskilled	−0.1
Participation Married Females		Unskilled (f), Skilled (m)	0.4
Unskilled	6.3	Skilled, Skilled	1.7
Skilled	2.0	Skilled (f), Unskilled (m)	0.6
Total	4.5	*All*	
		All Newborns	−3.2
		Winning Households	60.7

Note: Entries in the left panel show the effects (percentage changes) across steady states on selected aggregates, Entries in the right panel show the corresponding welfare effects (consumption compensation) for newborn households.

on average, while an equivalent single skilled female faces a loss of 1.2%. Nonetheless, since tax rates fall substantially, a majority of adults benefit from the elimination of the welfare state—about 60.7% of households benefit.

Overall, these findings highlight and anticipate trade-offs associated with reforming the welfare state. The welfare state targets transfers to low- and middle-income households. As a result, while they depress participation, hours, and output, they are highly valuable for some households and translate into substantial losses for all newborns associated with its elimination—even when tax rates are sharply reduced across the board. These losses mask gains for many agents, resulting in a significant majority of newborns in favor of this hypothetical move. The significant majority in favor of elimination of the system (60.7%) illustrates the trade-offs involved in an economy with substantial heterogeneity like ours.

6. RETHINKING THE WELFARE STATE

We now conduct several quantitative experiments in which we provide answers to the questions that motivate the paper. In all experiments, the rate of return of capital is constant across steady states—but rental prices for labor services change in order to be consistent with equilibrium conditions. We first consider replacing the current transfer scheme with a Universal Basic Income scheme (UBI) and then with a Negative Income Tax (NIT) across steady-state equilibria. We then discuss the effects of these reform cases when transitional dynamics are taken into account. In the benchmark economy, tax revenue is used to finance all the transfers and government spending G. When we eliminate existing transfers, we assume that the government adjusts taxes to cover G and transfers associated with UBI or NIT.

6.1. *A Universal Basic Income*

In our first experiment, each household receives a transfer per household member (including children) in all dates and states. The current welfare state is abolished while the tax system is unchanged. We dub this experiment a *Universal Basic Income* scheme (UBI).

Specifically, we search across steady states for the level of the UBI transfer that maximizes the ex ante welfare of all newborns. We balance the budget by adjusting the 'level' parameter of the tax function (λ) proportionally.

Our findings are presented in Tables III and IV. We find that a per-person transfer of about 3.2% of mean household income maximizes the welfare of newborns. This corresponds to about \$3100 per person in 2019 dollars (\$12,400 for a married household with two children at home). To balance the budget, tax rates need to substantially increase; for a married household with two children at mean income, the average rate increases to 14.4% from 9.2% in the benchmark.[12] This occurs as at the welfare-maximizing level, the aggregate expenditure on transfers increases in a significant way relative to the benchmark: from 2.3% in the benchmark case to 5.9%. The UBI transfers, coupled with higher taxes, depress hours, participation, and output across steady states. Total hours and output decline by 0.9% and participation rates of unskilled and skilled married females decline by 4.4% and 1.9%, respectively. Since hours worked for those away from the margin of indifference do not change much relative to the benchmark, per-worker hours for females increase, as Table III shows.

Table IV illustrates the welfare consequences of the UBI policy across steady states. Even at the best policy, ex ante welfare for all newborns declines, with a compensating variation of −1.3%. But a majority of newborns, 53.2%, support a UBI program. As it was the case with eliminating the current welfare system, lifetime-poor households suffer under the UBI, since it does not fully replace the transfers they were receiving in the benchmark economy while they see their taxes go up. This contributes to an overall welfare loss. Unskilled single females experience a welfare loss at birth of about 2.5%, and skilled ones a loss of about 0.7%. On the other hand, unskilled married households are strong winners as Table IV demonstrates, with a welfare gain of about 2.1% for a married couple with two unskilled adult members. This reflects that some low-to-middle-income households, who did not receive much in terms of transfers in the benchmark economy, now get a generous UBI transfer contributing to generating majority support.

Transitional Dynamics. What is the role of transitional dynamics in our welfare findings? To address this question, we compute the transitions associated to the non-anticipated introduction of the welfare-maximizing transfer discussed before. The budget is balanced by adjusting the level parameter of the tax function in each period of the transition across steady states.

Table IV shows that the welfare losses are lower for newborn households at the start of the transition than in the steady state—0.5% versus 1.3%. As the economy gradually shrinks along the transition, taxes to finance the UBI transfer increase, and therefore, welfare losses are larger when steady states are compared. Nonetheless, there is a majority of newborns who lose from the implementation of a UBI; the fraction of newborns who support the policy drops from 53.2% to 44.1%. If we only count all households alive at the start of the transition, there is still a majority against the implementation.

A UBI on Top of the Welfare State?. It is worth noting that the welfare-maximizing transfer level is substantially below the magnitudes advocated in policy discussions. For instance, there are proposals of a UBI transfer of \$1000 per month per adult, which is a much higher transfer than what we find as optimal. Indeed, we find that with transfer levels higher than the optimal one, welfare losses non-trivially increase and popular support dissipates quickly,

[12]The λ values in Table SA-VI are multiplied by 0.9433 to balance the government's budget.

A natural next question is: what if the UBI transfer is given on top of the existing welfare state? We find that a relatively small transfer of about 0.5% of mean household income maximizes ex ante welfare. There is a welfare gain of 0.1%, but only 48.9% of households are in favor of such a program. Larger transfers lead to welfare losses. For instance, if the transfer is 1% of the mean household income per person, an ex ante welfare loss emerges and 52.5% of adults oppose it. If, instead, we impose the transfer that maximizes welfare under the UBI reform (a transfer of 3.2% of mean income) on top of the existing programs, output losses are more significant, and ex ante welfare and popular support decline much further.

These findings follow from the fact that the welfare system is resilient, as its elimination leads to large welfare losses. Likewise, a common transfer to all is quite expensive and requires non-trivial tax hikes that depress welfare. We conclude from these findings that only small transfers on top of the existing welfare state—far from those advocated by proponents of a UBI—can marginally improve welfare but without majority support.

Summary. Overall, our findings indicate that a UBI policy reform is hard to justify on ex ante welfare grounds as a replacement for the current welfare state. Yet, it is supported by a majority despite its macroeconomic magnitude and the additional tax revenue it requires. A UBI policy on top of the current welfare state is not a good idea; only marginal welfare gains emerge for a small transfer and there is a clear majority against the move.

6.2. *A Negative Income Tax*

We now evaluate a more drastic reform that eliminates the current welfare state *and* the progressive income taxation. Specifically, we introduce a proportional income tax combined with a transfer for all, adult and children. Following Friedman (1962) and the literature that followed, we dub this linear income tax a *Negative Income Tax* system, or NIT for short. We again search for the welfare-maximizing per-household-member transfer and balance the budget by adjusting the proportional tax rate that applies to all households.

Table III shows the effects on aggregates. The transfer at the welfare-maximizing level is about 4.8% of mean household income, or \$4700 in 2019 dollars (\$18,800 for a married couple with two children). Thus, the welfare-maximizing NIT transfer is significantly more generous than the best one in the UBI case, and involves a drastic increase in resources devoted to redistribution—about 8.8% of output. The proportional tax rate that supports the welfare-maximizing NIT is 19.8%. Thus, tax rates are higher for most households—a married household with two children at around mean income faces an average tax rate of about 9.2% in the benchmark economy—but marginal rates are lower for those with top incomes. For instance, the marginal rate for a married household at three times mean income amounts to 21.4% in the benchmark case.

Overall, in net terms, the NIT reform leads to depressing effects on hours worked and output, as Table III demonstrates. Indeed, given its generosity at the welfare-maximizing case, its effects on total hours worked and participation rates are stronger than in the UBI case. But output declines less with the NIT than it does with the UBI (0.9% vs. 0.6%) since a proportional tax encourages higher labor supply from more productive households.

Welfare. Table IV shows that an NIT generates ex ante welfare gains of about 0.2% of consumption, which are accompanied by substantial majority support for the reform among newborns—more than two-thirds of newborns support the reform at birth. There

TABLE III

AGGREGATE AND INEQUALITY FINDINGS (CHANGES RELATIVE TO BENCHMARK).

	UBI Maximum Welfare	NIT Maximum Welfare	NIT (2) Maximum Welfare
Output	−0.9	−0.6	−0.4
Aggregate Hours	−0.9	−1.2	−1.0
Hours per worker (All Females)	0.8	0.5	−0.0
Hours per worker (All Males)	−0.5	−0.5	−0.2
Participation Married Females:			
Unskilled	−4.4	−6.0	−4.2
Skilled	−1.9	−2.3	−1.5
Total	−3.4	−4.4	−3.1
Proportional Tax Rate (%)	–	19.8	19.8
Transfer (% Household Income)	3.2	4.8	7.0, 4.1
Transfers (% Output)	5.9	8.8	8.7
Inequality			
△ Var Log Household Income	−0.027	0.01	0.009
△ Var Log Household Income (after taxes and transfers)	−0.006	0.029	−0.014

Note: Entries in the top panel show effects (percentage changes) across steady states on selected variables driven by the different quantitative experiments. Entries in the bottom panel are changes in the variance of log household income before and after tax and transfers, relative to the benchmark economy.

are of course winner and losers. Married households enjoy substantial welfare gains, which is most significant among couples with two unskilled partners. On the other hand, single female households as a group experience ex ante losses. The losses are more significant among unskilled and those with children.

In the Supplemental Material (see Figure SA10), we show how aggregate output, ex ante aggregate steady-state welfare, and majority support change for different levels of the NIT transfers. When the transfer equals zero, that is, when the tax system is simply a proportional tax with no transfers whatsoever, output is about 3.2% higher than in the benchmark case. But there is a substantial welfare cost. As transfers increase, tax rates, welfare, and popular support increase as well, but output declines. Altogether, as the lump-sum transfer increases, both welfare and support for the reform first sharply increase and then decline when the decline in output dominates.

If we consider transitional dynamics, the welfare gains associated to the welfare-maximizing transfer among newborns at the start of the transition are higher than under the comparisons across steady states (0.4% vs. 0.2%). A majority of newborn households support the policy shift (55.1%). This majority is even higher (59.4%) if we consider all households alive at the start of the transition.

Comparison With a Proportional Income Tax. Since the NIT has two elements, that is, the replacement of all transfers with a common payment to individuals and a move to a proportional tax system, it is informative to study what happens with a simple proportional tax that leaves the welfare state in place. We find that in this case, aggregate hours and output increase by about 1.6% and 1.4%, respectively, requiring a supporting tax rate of 10.8%. In terms of welfare, we find effectively no gains or losses on an ex ante basis. Again there are sharp differences between winners and losers. Our results show that skilled married couples have a welfare gain of 1.1%, whereas unskilled single females experience a loss of about 0.4%. A strong majority of newborn households, about 62.3% of them, are *against* it. Hence, generous transfers, made possible by a proportional income tax, are key for the success of NIT.

Differentiated Transfers. Since welfare gains from an NIT reform are unevenly distributed between married and single households and relatively small in the aggregate, we consider an NIT regime with transfers differentiated by the marital status of adults but with a common tax rate. We refer to this case as NIT(2). Specifically, we search for a transfer and ratio of transfers to individuals in married households relative to single households that maximize ex ante welfare *and* preserve majority support.

The welfare-maximizing transfer per person in single households is about 7% of the mean household income, while it is about 4.1% in married households (about $6860 and about $4000 in 2019 dollars). The tax rate that supports this arrangement is about the same as in the baseline NIT exercise (19.8%). Output and aggregate hours decline by 0.4% and 1.0% relative to the benchmark economy. This reform effectively means that a single female with two children, under an income level of one-half mean household income, would receive a net transfer after taxes of about 11.1% of mean household income (about $10,850 in 2019 dollars). The net transfer for a married couple with the same income and two children would be about 6.50% of mean household income (about $6350 in 2019 dollars).

An NIT reform of this type leads to ex ante welfare gain of about 0.7%, with 51.4% of adults supporting the reform in the welfare-maximizing scenario. Clearly, gains are larger than in the undifferentiated NIT. As earlier, if we consider transitions across steady states, welfare gains for newborn households are larger, and a substantial majority supports the policy move.

Inequality. The alternative tax-transfer schemes have ambiguous effects on inequality, as Table III shows. The variances of log household income, before and after taxes and transfers, decline with a shift to a UBI. Meanwhile, both measures increase with the NIT case, and move in opposite directions in the case of an NIT with differentiated transfers for marital status.

In understanding these results, it is key to bear in mind that several forces are at play here. First, the removal of transfers in the benchmark and the replacement for a smaller UBI transfer lead to increases in labor supply for some groups (e.g., low-skilled single females), contributing to a reduction in inequality in the UBI case. Second, an NIT reform involves a larger transfer alongside the full elimination of increasing marginal rates on household income. The result is a net increase in household income inequality. The same forces with differentiated transfers dictate larger increases in inequality before taxes and transfers in the second NIT experiment. Overall, the effects on inequality of drastically different alternatives appear of second order.[13]

Summary. Quantitatively, welfare gains under an NIT regime are substantially larger than the (negative) gains under a standard UBI scheme, and even larger gains can be obtained when NIT transfers are differentiated by marital status. Taking into account transitional dynamics makes these welfare findings even stronger. Overall, what accounts

[13]The changes in the log variances are quite small relative to the changes in variances due to taxes and transfers. For instance, while a baseline NIT increases the pre-tax and transfer inequality by about 1 log point, taxes and transfers reduce the variance by about 14.7 points.

for the relative success of a Negative Income Tax in terms of ex ante welfare and majority support? The upshot is that a larger degree of redistribution is feasible given the smaller tax distortions under an NIT regime, that is, with the elimination of increasing marginal tax rates and lower taxes on secondary earners. As tax distortions are reduced with a proportional tax, compared to the UBI case, the size of the aggregate economy is larger and collecting the tax revenues that are necessary to finance transfers becomes easier. The net result is that a higher transfer level becomes feasible under an NIT relative to a UBI scheme. Put differently, a drastic tax reform that reduces marginal tax rates at top incomes while making them common across all earners makes more extensive redistribution possible.

7. FINDINGS IN PERSPECTIVE

In this section, we place our results in perspective, with a focus on the effects and consequences of an NIT. We study an NIT reform in an economy with a lower level of inequality than our benchmark. We then evaluate an NIT reform when life-cycle facts are parameterized under a cohort instead of time effects. We also assess the use of progressive taxation to finance transfers and the effects of changes in the parameterization of preferences.

7.1. *Rethinking the Welfare State When Inequality Is Lower*

The U.S. economy has changed in critical ways relative to what it was in the recent past. In particular, there have been drastic changes in the demographic composition of U.S. households and in dimensions of inequality. For instance, only about 19% of females were skilled under our definition in 1980, while this figure more than doubled to nearly 39% in 2008, our baseline year for demographics. Substantial changes in marital sorting accompanied these changes; about 14% of married households were of the skilled-skilled category in 1980, while the corresponding figure in our parameterization is nearly 27%. Meanwhile, the skill premium was 1.4 under our assumptions in 1980, and it increased to 1.8 in our benchmark parameterization.

Since these changes can affect the importance of the redistributive component of the NIT vis-à-vis its effects on tax distortions, they could be of importance for our conclusions. To what extent do these underlying differences matter for the implications of an NIT? To address this question, we parameterize our economy to demographic and endowment targets of the past (circa 1980). We impose the demographics prevailing in 1980, use corresponding wage inequality over the life cycle to parameterize endowments, and force our economy to be consistent with the skill premium in 1980. We refer to this case as the *1980 economy*. We then search for the welfare-maximizing NIT as in our baseline exercises.[14]

We conduct our exercises in two levels. In our first case, we do not impose the participation rates of 1980 to parameterize the 1980 economy, while we do so in our second case. We summarize our results in Table V. We find that the welfare-maximizing NIT does not lead to a contraction in output in the long run. Furthermore, the resulting welfare gains are larger than in our baseline case (0.2% vs. 0.7–0.8%), with more substantial majority support, under an NIT transfer that is smaller—about 3% of household income. In short, these findings follow from the factors leading to lower underlying inequality among

[14]We describe the 1980 data in detail in the Supplemental Material.

TABLE IV

WELFARE EFFECTS (CONSUMPTION COMPENSATION, NEWBORNS, %).

	UBI Welfare	NIT Welfare	NIT (2) Welfare
Single F			
Unskilled	−2.5	−2.0	0.5
Skilled	−0.7	−0.7	0.3
No Child	−0.6	−0.6	0.3
Early Childbearing	−2.2	−1.7	0.3
Late Childbearing	−0.5	−0.4	0.2
Married			
Unskilled, Unskilled	2.1	2.8	0.2
Unskilled (f), Skilled (m)	0.2	0.3	−0.1
Skilled, Skilled	0.0	0.6	−0.3
Skilled (f), Unskilled (m)	0.2	0.4	−0.1
No Child	0.1	0.1	−0.2
Early Childbearing	1.4	2.3	−0.1
Late Childbearing	1.0	1.7	−0.0
All Newborns	−1.3	0.2	0.7
Winning Households	53.2	68.2	51.4
Including Transitional Dynamics			
All Newborns	−0.5	0.4	1.0
Wining Households	44.1	55.1	56.8
Wining Households (All Alive)	46.8	59.4	75.3

Note: Entries show the welfare effects (consumption compensation) driven by the reform of the welfare state. The top panel reports welfare gains across steady states with a constant interest rate across steady states. The calculations in the panel below take into account transitional dynamics between steady states.

households. These factors determine that an NIT reform requires a lower transfer to maximize ex ante welfare and a concomitant lower tax rate than in the baseline case. Hence, the detrimental effects on aggregates from taxes and transfers are of smaller magnitude, resulting in turn in the expansion (not contraction) of output and hours that Table V shows.

We conclude from these findings that an NIT reform in an economy with the characteristics of the United States circa 1980 is more appealing on several fronts, regardless of whether we consider changes in participation rates or not (case I or case II). Different factors leading to lower levels of inequality result in an environment more favorable to the implementation of an NIT reform.

7.2. *Cohort Effects in Data*

As described in Sections 3.1 and 4, our analysis relies on a parameterization in which life-cycle facts were characterized controlling by time effects. In the Supplemental Material, we describe our main empirical findings when we extract age profiles controlling instead by cohort effects. Under cohort effects, we find steeper profiles of hourly wages as well as steeper profiles of wage inequality relative to the time-effects case. These findings are in line with the literature.

What are the implications of assuming cohort effects in our parameterization for our results? Table V summarizes our findings. We find a welfare-maximizing NIT transfer of around 4.7% of mean household income, roughly at the same level of the baseline case under time effects. Welfare gains are higher, at about 0.8% for all newborn households.

Aggregate output and hours expand across steady states by 0.7% and 0.4%, respectively, unlike the baseline case under time effects.

Qualitatively, these results are unsurprising. Steeper inequality profiles lead to higher variances of persistent shocks relative to the time-effects case, and therefore, there is a larger role of transfers for all provided by an NIT transfer. Likewise, steeper wage profiles lead to potentially larger gains related to the replacement of increasing marginal rates of household taxation. Hence, larger welfare gains are available. We conclude from these findings that parameterizing our economy under a cohort-effects view of the life-cycle data makes a basic NIT reform more desirable.

7.3. *Funding Transfers With Higher Progressivity*

In our main exercises, common transfers are funded by taxing household income in the least distorting ways possible: either by shifting the tax function (UBI) or by using a proportional income tax (NIT). What if transfers are financed by changing the progressivity of the tax function? Can this improve upon the UBI or NIT schemes analyzed earlier?

To address this question, we proceed as follows. We eliminate existing transfers, and search for the level of tax progressivity (curvature) of the income tax function and a common transfer to all individuals that are consistent with budget balance and maximize ex ante welfare. Since the curvature of tax functions differs (married, single, etc.), we search for the common factor that increases or reduces the level of progressivity in all cases. Note here that increasing the curvature of the tax function implies a rotation in the tax function; it increases average rates at the top, but reduces average tax rates at the bottom.

We find a welfare-maximizing transfer relatively small, of only about 1.25% of mean household income, and that this transfer level implies a non-trivial *reduction* in the curvature of the tax functions—about 60% reduction across the board. This determines an ex ante welfare loss of about -1.8%. Put differently, the welfare-maximizing transfer-progressivity combination does not dominate the status quo and leads to a worse outcome than our basic UBI.

To understand these findings, it is important to note that increases in progressivity from the benchmark levels, as they lead to steeper marginal rates, lead to non-trivial declines in economic aggregates and can only increase revenue marginally in the long run. Guner et al. (2016) made explicitly this point in the context of life-cycle model with heterogeneity. Using revenues to provide for transfers to all makes the problem of obtaining additional revenues even worse. We conclude from these findings that a shift to common transfers to all hinges critically on imposing taxes that distort decisions little, either by shifting taxes for all (UBI) or by drastically changing the nature of income taxation (NIT).

7.4. *CRRA Preferences and Economies of Scale*

We have conducted all of our analysis using preferences in which preferences for consumption are logarithmic, with a coefficient of relative risk aversion of 1. Likewise, our analysis did not include scale economies in consumption as scale economies are not relevant with log preferences, but scale economies could become relevant with CRRA preferences. What are the implications for our conclusions regarding an NIT reform if CRRA preferences and scale economies are considered?

To answer this question, we set the relative risk aversion to a higher level (1.5) and impose a commonly used adjustment for scale economies, equal to the square root of the number of people in the household. We then parameterize the model economy again in

TABLE V

FINDINGS IN PERSPECTIVE (% CHANGES RELATIVE TO BENCHMARK).

	Baseline Findings	1980 (I)	1980 (II)	Cohort Effects	CRRA & Scale Economies
Output	−0.6	0.7	0.6	0.7	−3.8
Aggregate Hours	−1.2	0.9	0.8	0.4	−5.8
Participation Married Females:					
Unskilled	−6.0	−0.1	−0.3	−1.1	−18.3
Skilled	−2.3	0.3	−0.2	1.4	−8.1
Total	−4.4	0.1	−0.3	0.0	−13.9
Welfare					
All Newborns (%)	0.2	0.8	0.7	0.8	0.8
Winning Households	68.2	80.0	81.0	67.5	49.5
Proportional Tax Rate (%)	19.8	14.8	14.8	19.6	30.8
Transfer (% Household Income)	4.8	3.0	3.0	4.7	8.2
Transfers (% Output)	8.8	5.9	5.9	8.6	15.6

Note: Entries in the top panel show effects of a welfare-maximizing NIT experiment. The second column shows the baseline findings for comparison. The third column displays the results when shocks and demographics are calibrated to data circa 1980, but married female participation rates are from the benchmark. The fourth column shows the corresponding results when shocks, demographics, and participation rates are calibrated to data from circa 1980. The fifth column shows the results when the benchmark economy is calibrated under a cohort-effects view of the data. The final column shows results when preferences allow for a risk aversion coefficient exceeding 1 and scales economies in household consumption.

this scenario and find the welfare-maximizing NIT arrangement. Our findings are summarized in Table V. In this scenario, an NIT experiment leads to much larger transfers and associated tax rates, and to substantial declines in output and labor supply aggregates. The welfare-maximizing transfers to all reach 8.2% with a tax rate of about 31%. Welfare gains are of about 0.8%, but without a majority of households in support of the NIT reform.[15]

In interpreting these results, it is important to note that this exercise is a major departure in our analysis and should be taken with caution. Our benchmark parameterization is consistent with a host of facts, including the growth of consumption dispersion over the life cycle, and in macroeconomic terms, it is consistent with balanced growth. This alternative case is not consistent with balanced growth, with income effects dominating substitution effects. This implies that increasing wages for all would reduce labor supply and that increasing tax rates would *increase* it. We would then expect that much higher transfers emerge in the welfare-maximizing case, with concomitantly large depressing effects on aggregates, as our results show.

8. CONCLUDING REMARKS

Three main points emerge from our analysis so far. First, it is hard to improve upon the current structure of the welfare state via simple transfer schemes. Transfers to poorer households are highly valued, and thus, any reform to the current system needs to confront the fact that non-trivial resources accrue to poorer households. As a result, simpler

[15] If we impose a CRRA coefficient of 1.5 without any scale economies in consumption, calibrating the economy to data, the results are more moderate but similar. Output and hours worked decline by 3.1 and 4.9% across steady states, respectively, under a transfer of about 7.75% of mean household income. Welfare gains amount to 0.4%.

schemes that maximize ex ante welfare relative to the status quo require drastic changes in taxation.

Second, a UBI scheme is generically not a good idea and is dominated by an NIT. Why? Considerable resources need to be transferred to poorer households for their welfare not to fall. And since transfers would accrue to all individuals, taxes need to increase substantially, leading to ex ante welfare losses. A generous UBI transfer imposed on top of the current welfare state, or a UBI transfer financed via increases in tax progressivity, do not change this conclusion. It follows that in our economy, a UBI scheme, as proposed in popular discussions, is not a good idea.

Last, NIT arrangements generate ex ante welfare gains and lead to popular support due to the associated reduction or elimination of pre-existing distortions (i.e., increasing marginal tax rates, household rather than individual income-tax base) and the concomitant increase in output and revenues. Our findings hold when transitional dynamics are taken into account and when life-cycle data are interpreted via cohort effects instead of time effects. Interestingly, we also find that the desirability of an NIT reform is higher for an economy with the characteristics of the U.S. economy in the past (circa 1980).

We end this paper with three comments. First, the administrative costs of running a welfare state can be large. Isaacs (2008), for example, calculated that the costs of running Food Stamps, Housing Subsidies, and the TANF programs are as high as 15 cents per each dollar benefit issued. Our analysis abstracts from such administrative costs, and hence might underestimate the potential benefits of moving to a simpler system like the NIT. Second, a variant to the NIT system could use a consumption tax instead of a flat-rate income tax. As a consumption tax does not distort capital formation, this implementation could lead to larger gains in output, labor supply, and welfare than we found in our analysis. Finally, it might be important to evaluate reforms to the welfare state when health and health-related transfers are taken into account. In this vein, we consider that a model of endogenous health, where both medical and non-medical means-tested transfers are modeled and non-medical income support can affect investments in health, is a promising line of future research.[16] We leave this and other issues for future research.

REFERENCES

AGUIAR, MARK, AND ERIK HURST (2013): "Deconstructing Life Cycle Expenditure," *Journal of Political Economy*, 121, 437–492. [2267]

AIYAGARI, S. RAO, JEREMY GREENWOOD, AND NEZIH GUNER (2000): "On the State of the Union," *Journal of Political Economy*, 108, 213–244. [2275]

ATTANASIO, ORAZIO P., JAMES BANKS, COSTAS MEGHIR, AND GUGLIELMO WEBER (1999): "Humps and Bumps in Lifetime Consumption," *Journal of Business and Economic Statistics*, 17, 22–35. [2272]

ATTANASIO, ORAZIO P., HAMISH LOW, AND VIRGINIA SÁNCHEZ-MARCOS (2005): "Female Labor Supply as Insurance Against Idiosyncratic Risk," *Journal of the European Economic Association*, 3, 755–764. [2272]

BADEL, ALEJANDRO, MARK HUGGETT, AND WENLAN LUO (2020): "Taxing Top Earners: A Human Capital Perspective," *The Economic Journal*, 130, 1200–1225. [2264]

BAYER, CHRISTIAN, AND MORITZ KUHN (2019): "Which Ladder to Climb? Decomposing Life-Cycle Wage Dynamics," IZA Discussion Paper No. 12473. [2266]

BITLER, MARIANNE, JONATHAN GELBACH, HILARY HOYNES, AND MADELINE ZAVODNY (2004): "The Impact of Welfare Reform on Marriage and Divorce," *Demography*, 41, 213–236. [2275]

BLUNDELL, RICHARD, MONICA COSTA DIAS, COSTAS MEGHIR, AND JONATHAN SHAW (2016): "Female Labor Supply, Human Capital and Welfare Reform," *Econometrica*, 84, 1705–1753. [2276]

[16]Mahler and Yum (2022) for instance, built a model of endogenous health where changes in healthy behavior are subject to adjustment costs.

BLUNDELL, RICHARD, LUIGI PISTAFERRI, AND ITAY SAPORTA-EKSTEN (2016): "Consumption Inequality and Family Labor Supply," *American Economic Review*, 106, 387–435. [2264]
BOAR, CORINA, AND VIRGILIU MIDRIGAN (2022): "Efficient Redistribution," *Journal of Monetary Economics*, 131, 78–91. [2264]
BORELLA, MARGHERITA, MARIACRISTINA DE NARDI, AND FANG YANG (2023): "Are Marriage-Related Taxes and Social Security Benefits Holding Back Female Labour Supply?" *Review of Economic Studies*, 90, 102–131. [2264]
CONESA, JUAN CARLOS, LI BO, AND QIAN LI (2023): "A Quantitative Evaluation of Universal Basic Income," *Journal of Public Economics*, 223, 104881. [2264]
DARUICH, DIEGO, AND RAQUEL FERNANDEZ (2023): "Universal Basic Income: A Dynamic Assessment," NBER Working Paper 27351. [2264]
FERRIERE, AXELLE, PHILIPP GRUBENER, GASTON NAVARRO, AND OLIKO VARDISHVILI (2023): "On the Optimal Design of Transfers and Income-Tax Progressivity," *Journal of Political Economy Macroeconomics*, 1 (2), 276–333. [2264]
FRIEDMAN, MILTON (1962): *Capitalism and Freedom*. Chicago, London: University of Chicago Press. [2285]
GOLDIN, CLAUDIA (2021): *Career and Family: Women's Century-Long Journey Towards Equality*. Princeton, NJ: Princeton University Press. [2266]
GREENWOOD, JEREMY, NEZIH GUNER, GEORGI KOCHARKOV, AND CEZAR SANTOS (2016): "Technology and the Changing Family: A Unified Model of Marriage, Divorce, Educational Attainment, and Married Female Labor-Force Participation," *American Economic Journal: Macroeconomics*, 8, 1–41. [2275]
GUNER, NEZIH, REMZI KAYGUSUZ, AND GUSTAVO VENTURA (2012): "Taxation and Household Labour Supply," *Review of Economic Studies*, 79, 1113–1149. [2264]
——— (2020): "Child-Related Transfers, Household Labor Supply and Welfare," *Review of Economic Studies*, 87, 2290–2321. [2264,2277]
——— (2023)a: "Supplement to 'Rethinking the Welfare State'," *Econometrica Supplemental Material*, 91, https://doi.org/10.3982/ECTA19921. [2265]
——— (2023)b: "Rethinking the Welfare State," CEMFI Working Paper No. 2304. [2272,2277,2282]
GUNER, NEZIH, MARTIN LOPEZ-DANERI, AND GUSTAVO VENTURA (2016): "Heterogeneity and Government Revenues: Higher Taxes at the Top?" *Journal of Monetary Economics*, 80, 69–85. [2264,2290]
GUVENEN, FATIH (2009): "An Empirical Investigation of Labor Income Processes," *Review of Economic Dynamics*, 12, 58–79. [2274]
HEATHCOTE, JONATHAN, KJETIL STORESLETTEN, AND GIANLUCA VIOLANTE (2005): "Two Views of Inequality Over the Life Cycle," *Journal of the European Economic Association*, 3, 765–775. [2266]
——— (2010): "The Macroeconomic Implications of Rising Wage Inequality in the United States," *Journal of Political Economy*, 118, 681–722. [2277]
——— (2017): "Optimal Tax Progressivity: An Analytical Framework," *Quarterly Journal of Economics*, 132, 1693–1754. [2264]
HOLTER, HANS A., DIRK KRUEGER, AND SERHIY STEPANCHUK (2019): "How Do Tax Progressivity and Household Heterogeneity Affect Laffer Curves?" *Quantitative Economics*, 10, 1317–1356. [2264]
HOSSEINI, ROOZBEH, KAREN A. KOPECKY, AND KAI ZHAO (2021): "How Important Is Health Inequality for Lifetime Earnings Inequality?" FRB Atlanta Working Paper No. 2021-1. [2276]
HOYNES, HILARY, AND JESSE ROTHSTEIN (2019): "Universal Basic Income in the US and Advanced Countries," *Annual Review of Economics*, 11, 929–958. [2264]
HUGGETT, MARK, GUSTAVO VENTURA, AND AMIR YARON (2011): "Sources of Lifetime Inequality," *American Economic Review*, 101, 2923–2954. [2274,2277,2279]
ISAACS, JULIA (2008): "The Costs of Benefit Delivery in the Food Stamp Program. United States Department of Agriculture," Contractor and Cooperator Report No. 39. [2292]
JAIMOVICH, NIR, ITAY SAPORTA-EKSTEN OFFER SETTY, and YANIV YEDID-LEVI (2022): "Universal Basic Income: Inspecting the Mechanisms," IZA Discussion Paper No. 15058. [2264]
KATZ, LAWRENCE, AND KEVIN MURPHY (1992): "Changes in Relative Wages, 1963-1987: Supply and Demand Factors," *The Quarterly Journal of Economics*, 107, 35–78. [2279]
KAYGUSUZ, REMZI (2010): "Taxes and Female Labor Supply," *Review of Economic Dynamics*, 13, 725–741. [2264]
——— (2015): "Social Security and Two-Earner Households," *Journal of Economic Dynamics and Control*, 59, 163–178. [2264]
KEARNEY, MELISSA S. (2004): "Is There an Effect of Incremental Welfare Benefits on Fertility Behavior? A Look at the Family Cap," *Journal of Human Resources*, 39, 295–325. [2275]
KINDERMANN, FABIAN, AND DIRK KRUEGER (2020): "High Marginal Tax Rates on the Top 1%? Lessons From a Life Cycle Model With Idiosyncratic Income Risk," *American Economic Journal: Macroeconomics*, 14, 319–366. [2264]

KUHN, MORITZ, AND JOSÉ-VÍCTOR RIOS-RULL (2016): "2013 Update on the US Earnings, Income, and Wealth Distributional Facts: A View From Macroeconomics," *Federal Reserve Bank of Minneapolis Quarterly Review*, 37 (1), 2–73. [2279]

LOPEZ-DANERI, MARTIN (2016): "NIT Picking: The Macroeconomic Effects of a Negative Income Tax," *Journal of Economic Dynamics and Control*, 68, 1–16. [2264]

LOW, HAMISH, COSTAS MEGHIR, AND LUIGI PISTAFERRI (2022): "Marriage, Labor Supply and the Dynamics of the Social Safety Net," NBER Working Paper No. 24356. [2264]

LUDUVICE, ANDRE VICTOR DOHERTYRE (2021): "The Macroeconomic Effects of Universal Basic Income Programs," Federal Reserve Bank of Cleveland Working Paper No. 21-21. [2264]

MAHLER, LUKAS, AND MINCHUL YUM (2022): "Lifestyle Behaviors and Wealth-Health Gaps in Germany," CEPR Discussion Paper No. DP17036. [2292]

MOFFITT, ROBERT A. (2003): "The Negative Income Tax and the Evolution of U.S. Welfare Policy," *Journal of Economic Perspectives*, 17, 119–140. [2264,2277]

MOFFITT, ROBERT J., BRIAN J. PHELAN, AND ANNE E. WINKLER (2020): "Welfare Rules, Incentives, and Family Structure," *Journal of Human Resources*, 55, 1–42. [2275]

OECD (2019): *Social Expenditure Update 2019: Public Social Spending Is High in Many OECD Countries*. Paris: OECD Publishing. [2262]

ORTIGUEIRA, SALVADOR, AND NAWID SIASSI (2013): "How Important Is Intra-Household Risk Sharing for Savings and Labor Supply?" *Journal of Monetary Economics*, 60 (650), 666. [2264]

——— (2022): "The U.S. Tax-Transfer System and Low-Income Households: Savings, Labor Supply, and Household Formation," *Review of Economic Dynamics*, 44, 184–210. [2264]

PRIMICERI, GIORGIO, AND THIJS VAN RENS (2009): "Heterogeneous Life-Cycle Profiles, Income Risk and Consumption Inequality," *Journal of Monetary Economics*, 56, 20–39. [2267]

RAUH, CHRISTOPHER, AND MARCELO R. SANTOS (2022): "How Do Transfers and Universal Basic Income Impact the Labor Market and Inequality?" CEPR Discussion Paper No. DP16993. [2264]

RECTOR, ROBERT, AND VIJAY MENON (2018): "Understanding the Hidden $1.1 Trillion Welfare System and How to Reform It," Backgrounder No. 3284, Heritage Foundation. [2277]

VAN PARIJS, PHILIPPE, AND YANNICK VANDERBORGHT (2017): *Basic Income: A Radical Proposal for a Free Society and a Sane Economy*. Cambridge, MA: Harvard University Press. [2264]

VOENA, ALESSANDRA (2015): "Yours, Mine, and Ours: Do Divorce Laws Affect the Intertemporal Behavior of Married Couples?" *American Economic Review*, 105, 2295–2332. [2275]

WU, CHUNZAN, AND DIRK KRUEGER (2021): "Consumption Insurance Against Wage Risk: Family Labor Supply and Optimal Progressive Income Taxation," *American Economic Journal: Macroeconomics*, 13, 79–113. [2264]

Co-editor Charles I. Jones handled this manuscript.

Manuscript received 5 July, 2021; final version accepted 18 June, 2023; available online 31 July, 2023.