On the optimal design of transfers and income-tax progressivity*

Axelle Ferriere	Philipp Grübener	Gaston Navarro	Oliko Vardishvili
PSE & CNRS	Goethe University	Federal Reserve Board	UC Irvine

March 13, 2023

Abstract

We study the optimal design of means-tested transfers and progressive income taxes. In a simple analytical model, we show that adding a transfer to a log-linear tax induces welfare gains almost as large as in the second-best allocation. Transfers allow for more progressive average than marginal tax-and-transfer rates, achieving redistribution while preserving efficiency. In a rich dynamic model, we quantify the optimal fiscal plan. We use new flexible functions featuring targeted transfers and progressive income taxes, delivering a good empirical fit across the income distribution. Transfers should be larger than currently in the U.S. and financed with moderate income-tax progressivity.

Keywords: Heterogeneous Agents, Fiscal Policy, Optimal Taxation, Redistribution.
JEL: E21, E62, H21, H23, H53.

*axelle.ferriere@psemail.eu, gruebener@econ.uni-frankfurt.de, gaston.m.navarro@frb.gov, oliko.vardishvili@gmail.com.

This work was supported by computational resources provided by the BigTex High Performance Computing Group at the Federal Reserve Bank of Dallas and by the Goethe-HLR cluster at Goethe University Frankfurt. An earlier version of this paper circulated under the title "Larger transfers financed with more progressive taxes? On the optimal design of taxes and transfers." We thank Árpád Ábrahám, Isaac Baley, Anmol Bhandari, Moritz Drechsel-Grau, Dirk Krueger, Abdoulaye Ndiaye, Facundo Piguillem, Dominik Sachs, Kjetil Storesletten, Gianluca Violante, and workshop and seminar participants at various seminars and conferences for helpful comments and suggestions. The views expressed are those of the authors and not necessarily those of the Federal Reserve Board or the Federal Reserve System. Edited by Anmol Bhandari.

1 Introduction

High levels of inequality have made redistributive policies a core topic in recent policy debates. Two key components of redistributive policies are means-tested transfers and progressive income taxes. Both policies can significantly alter the income distribution. In the United States, targeted transfers amount to almost 25% of income for the poorest income quintile, while income taxes reduce richest-quintile income by about 30%.[1] The optimal design of these two tools, targeted transfers and progressive income taxes, is thus of paramount importance.

In this paper, we study the joint optimal design of targeted transfers and progressive income taxes. We do so in two steps. We first develop a simple analytical model to inspect the role of transfers in the optimal fiscal system, and characterize their interaction with income-tax progressivity. Second, we calibrate a rich dynamic model of the U.S. economy and use it to quantify the optimal levels of transfers and income-tax progressivity. We pursue a Ramsey approach and endow a planner with flexible functions for transfers and income taxes, which we refer to as the tax-and-transfer ($t\&T$) system.

Our approach has several advantages. The functional forms we use resemble current policies implemented by many countries. Thus, our analysis sheds light on the optimal use of currently available instruments. Furthermore, the Ramsey approach is suitable for a rich quantitative evaluation with empirically realistic efficiency and redistribution concerns. Finally, the instruments we use are simple and characterized by few economically intuitive parameters. Yet they are flexible enough to generate nonlinear, and potentially non-monotonic, overall $t\&T$ schedules, features often found to be optimal in the Mirrleesian literature. Thus, the flexibility of our Ramsey approach allows us to build new intuition on the typical optimal taxation tradeoffs.

We present two main findings. First, the optimal $t\&T$ system typically features more progressive average than marginal rates so as to redistribute while preserving efficiency. Adding a transfer to progressive taxes precisely allows the planner to disentangle average from marginal rates, thus generating large welfare gains—systematically close to the second-best allocation in the simple analytical model. Second, using the dynamic quantitative model, we find that transfers should be larger than currently in the United States and financed with moderate income-tax progressivity. We further use the quantitative model to analyze how the income distribution shapes the optimal $t\&T$ system and to explore the quantitative importance of non-monotonic marginal $t\&T$ rates on welfare.

The analytical model builds on the work in Heathcote, Storesletten, and Violante (2017). We assume a continuum of households who face a static consumption-labor supply decision subject to idiosyncratic labor risk. We endow a utilitarian planner with a log-linear income tax and a lump-sum transfer. The key role of transfers has long been

[1] See Section 3.4 and Appendix B.1 for definitions of income, taxes, and transfers.

emphasized in the Mirrlees literature (Saez 2001). Furthermore, as we show, adding a transfer to the log-linear tax function substantially improves the fit to the U.S. $t\&T$ system.

We use local approximations in the analytical model to derive a closed-form formula for welfare that shows an optimal negative relation between transfers and income-tax progressivity, due to both *efficiency* and *redistribution* concerns. Higher income-tax progressivity and larger transfers both discourage labor supply. In turn, a planner finances larger transfers with lower progressivity to restore labor supply incentives. We refer to this channel as the *efficiency* concern. In terms of *redistribution*, a planner aims to decrease dispersion in consumption. Higher income-tax progressivity and larger transfers both reduce consumption dispersion. The larger the tax progressivity, the lower the consumption dispersion and the smaller the welfare gain from transfers. Thus, larger transfers are optimally financed with less progressive income taxes.

We then use an illustrative calibration of this simple model to globally compute the optimal transfers and income-tax progressivity. We find that transfers should be large and, thus, income-tax progressivity should be low. Redistribution is achieved via generous transfers, and thus progressive average rates, while efficiency is preserved with low income-tax progressivity, and thus flatter marginal rates. Furthermore, the optimal log-linear tax with a transfer generates welfare gains almost as large as the Mirrlees allocation—a result that we show to be robust across several calibrations. Intuitively, the planner has two instruments, a transfer and an income-tax progressivity, to achieve two targets, redistribution and efficiency.

We confirm these findings in a rich, quantitative Bewley-Huggett-Aiyagari incomplete-market model calibrated to the U.S. economy. We enrich the set of fiscal instruments and endow the planner with targeted transfers that, in line with current practices, phase out with both labor and capital income. The phase-out of transfers, combined with progressive income taxes, allows for non-monotonic marginal $t\&T$ rates. We estimate the tax and the transfer functions using household-level data. We also incorporate an empirically realistic process for idiosyncratic labor risk, with a thick right tail of productivity and higher-order moments of income risk, to accurately capture redistribution needs in the economy. The model produces labor supply elasticities and wealth effects in line with the data and thus accurately captures efficiency concerns as well. We consider once-and-for-all changes to the fiscal instruments and incorporate transitions toward the new steady state in our welfare computations.

The optimal $t\&T$ system in the quantitative model is substantially more redistributive than the current system in the United States. Optimal transfers amount to $19,800 (2012 U.S. dollars) per year for the lowest-income household.[2] This number implies an income

[2]As we discuss in Section 3.4, we report transfers per household, as it is the unit of calibration of the quantitative model.

floor of 23% of mean income. Transfers optimally phase out, albeit at a slow rate, such that a household with median income receives a transfer of about \$3,700. In line with our analytical findings, large transfers are optimally financed with moderate income-tax progressivity, at a level only slightly higher than in the current U.S. system. The optimal system implies large welfare gains of 6.00% in consumption equivalent terms. These welfare gains are largely due to better insurance and redistribution. In addition, despite a large drop in aggregate labor, there are small efficiency improvements because of a better allocation of hours worked across households. While welfare gains primarily accrue to the poor, 76% of households benefit from the reform.

A substantial part of the welfare gains in the benchmark plan can be achieved with either of two common tax proposals: a universal basic income (UBI) or an affine plan. In terms of tax instruments, the UBI plan eliminates the phase-out of transfers but still optimizes the progressivity of income taxes. The affine plan is a UBI financed with flat income taxes. The optimal UBI plan implies a lump-sum transfer of \$18,700 per household, financed with barely progressive income taxes: the marginal income-tax rate averages 56% for the poorest income quintile and 61% for the richest income quintile. As the optimal UBI plan features almost flat taxes, the optimal affine plan resembles the UBI one. The affine plan features a lump-sum transfer of \$20,300 and a tax rate of 60%. The welfare gain of the UBI is 5.36% in consumption equivalent terms, while the affine gains are 5.26%, both close to the gains in the benchmark plan. Thus, our framework is supportive of a lump-sum transfer provided that it is financed with the right income-tax progressivity. Alternatively, a plan with transfers phasing out implies lower income-tax rates, which may be easier to implement in practice.

The UBI and affine experiments point to the importance of disentangling average and marginal rates and shed light on the relative importance of the two standard elements in the Mirrlees literature: the positive intercept and the U-shape of the marginal rates. Transfers, even if lump sum, allow to separate the progressivity of average and marginal $t\&T$ rates. In contrast, the optimal log-linear plan with no transfers tightly links marginal to average rates and thus delivers smaller welfare gains, at only 2.88%. Additionally, the phasing out of transfers allows for non-monotonic marginal rates, which delivers only modest additional welfare gains.

Finally, we use our quantitative model to show that the income distribution shapes the negative relation between transfers and income-tax progressivity. In particular, the right tail of the income distribution shifts this relation: the lower the income concentration at the top, the lower the income-tax progressivity for each level of transfer. Yet, income concentration at the top barely affects the optimal level of transfers, and only the income-tax progressivity adjusts. In contrast, when bottom-income households are richer, the optimal income-tax progressivity for each transfer does not change, but the optimal transfers decrease—implying larger income-tax progressivity, despite lower inequality. In brief, optimal transfers are driven by the left tail of the income distribution, whereas income-tax progressivity is determined by the right tail of the distribution.

Related literature. This paper belongs to a large literature on optimal taxation. Seminal papers in the Mirrleesian tradition include Mirrlees (1971), Diamond (1998), and Saez (2001). At a general level, our Ramsey approach with flexible $t\&T$ functional forms allows for a quantitative evaluation of the main tradeoffs discussed in this literature within a rich dynamic model.

Our paper relates more closely to Heathcote and Tsujiyama (2021), which compares a fully nonlinear Mirrlees plan to a log-linear tax system in a static environment. In a setup comparable to theirs, we show that adding an intercept to a log-linear income-tax function comes systematically close to the Mirrlees plan—a result that holds across several calibrations, including theirs (see Section 2). We further consider targeted transfers and embed our analysis in a dynamic general equilibrium model in Section 4. Also related to our paper are Chang and Park (2021) and Park (2022), which use a variational approach to derive a tax formula in an Aiyagari model with Greenwood, Hercowitz, and Huffman (1988) preferences. Our Ramsey approach allows for intuitive insights as well as the inclusion of transitions from the calibrated steady state in welfare computations.[3]

Our quantitative framework relates to several papers interested in optimal tax progressivity in incomplete-market models. An early contribution is Conesa and Krueger (2006), and subsequent work has focused on: transitional dynamics (Bakış, Kaymak, and Poschke 2015), superstars and entrepreneurs (Brüggemann 2021; Kindermann and Krueger 2022), human capital accumulation (Badel, Huggett, and Luo 2020; Krueger and Ludwig 2016; Peterman 2016), negative income tax programs (Lopez-Daneri 2016), and Laffer curves (Guner, Lopez-Daneri, and Ventura 2016; Holter, Krueger, and Stepanchuk 2019). Our paper also relates to recent quantitative analyses of UBI policies, with a specific focus on human capital (Daruich and Fernández 2022; Luduvice 2021), labor market frictions (Jaimovich, Saporta-Eksten, Setty, and Yedid-Levi 2022; Rauh and Santos 2022), and consumption taxes (Conesa, Li, and Li 2021).

Two recent works analyze transfers and progressive taxes in a rich quantitative setup. Guner, Kaygusuz, and Ventura (2021) develops a rich modeling of the household, while Boar and Midrigan (2022) focuses more specifically on the role of wealth taxes. In line with our quantitative finding, both papers find that an affine plan is welfare improving.[4] Our Ramsey approach with flexible tax instruments allows for new insights on the

[3] Our paper also relates to advances in New Dynamic Public Finance—such as Kapička (2013); Farhi and Werning (2013); Golosov, Troshkin, and Tsyvinski (2016); Findeisen and Sachs (2017); Stantcheva (2017); and Boerma and McGrattan (2020)—and connects with the literature using microsimulations to evaluate fiscal policies; see Colombino and Islam (2022) for an example evaluating Universal Basic Income proposals.

[4] See also the recent work in Carroll, Luduvice, and Young (2022) for similar findings.

tradeoff between transfers and income-tax progressivity. Additionally, the phase-out of transfers allows us to consider non-monotonic marginal rates, a common property of the optimal Mirrleesian $t\&T$ system.

Roadmap. Section 2 presents the analytical model. Section 3 contains the quantitative model and its calibration. Section 4 analyzes the optimal fiscal plan. Section 5 discusses the effect of income distribution on optimal transfers and income-tax progressivity and other robustness checks. Section 6 concludes.

2 A simple analytical model

We start with a simple analytical model to highlight the key role of transfers in the optimal fiscal system. We build on Heathcote, Storesletten, and Violante (2017) and propose a tractable heterogeneous-agent model in which a government uses a log-linear income tax to finance public spending and a lump-sum transfer. We use local approximations to derive an optimal negative relation between transfers and income-tax progressivity, and further show that transfers should typically be positive. We resort to numerical methods and obtain large optimal transfers in this simple economy. In addition, welfare gains under the optimal plan are almost as large as in the second-best (Mirrlees) allocation.

We end this section showing that the empirical fit to the U.S. $t\&T$ system substantially improves when adding an intercept to the log-linear tax. Importantly, the fit improves the most for bottom and top income earners, precisely where the tax matters for redistribution and efficiency purposes.

2.1 Environment: A static economy

The economy is populated by a continuum of ex-ante homogeneous households, a representative firm, and a utilitarian government. Households are hand-to-mouth, they value consumption c and leisure $1 - n$, and their labor productivity z is log-normally distributed. The representative firm uses a linear technology to transform labor into output. The government finances exogenous government spending G and a lump-sum transfer T with log-linear labor taxes.[5]

Taxes.—A household with labor income y pays taxes $\mathcal{T}(y) = y - \lambda y^{1-\tau}$, where τ captures the progressivity and λ the level of taxes. For $\tau = 0$, tax rates are flat and equal to $1-\lambda$. When $\tau > 0$ ($\tau < 0$), marginal and average tax rates are increasing (decreasing) in income. When $\tau = 1$, after-tax income is equalized $\forall y$, implying full redistribution.

[5]This tax function has been widely used since Feldstein (1969), and more recently by Benabou (2000) and Heathcote, Storesletten, and Violante (2017).

Households.—Household i chooses consumption c_i and labor n_i to maximize utility

$$u(c_i, n_i) = \ln c_i - B\frac{n_i^{1+\varphi}}{1+\varphi} \tag{1}$$

subject to a static budget constraint

$$c_i = \lambda \left(z_i n_i\right)^{1-\tau} + T. \tag{2}$$

We assume a log-normal distribution for idiosyncratic productivity z, with variance v_ω controlling the degree of heterogeneity across households. The log-normal assumption, together with log utility and hand-to-mouth households, is convenient to derive a closed-form solution for welfare when transfers are zero.[6]

Government.— The government budget constraint is

$$G + T = \int z_i n_i di - \lambda \int \left(z_i n_i\right)^{1-\tau} di. \tag{3}$$

Technology.— The resource constraint is

$$\int c_i di + G = \int z_i n_i di. \tag{4}$$

Next, we derive a formula for welfare as a function of progressivity τ in the tractable case of $T = 0$. Then, we use local approximations around that point to characterize welfare as a function of τ and T. All derivations in this section are relegated to Appendix A.

2.2 Optimal income-tax progressivity with no transfers

When transfers are zero, the optimal labor policy is constant across households. Maximizing the household's utility (1) given the budget constraint (2) yields an optimal labor policy function in closed form,

$$n_0(\tau) \equiv \left(\frac{1-\tau}{B}\right)^{\frac{1}{1+\varphi}}, \tag{5}$$

where the 0 subscript stands for zero transfers. Equation (5) shows a one-to-one negative relationship between n and τ. Given this simple policy function, we can compute output,

[6]Without transfers, a no-trade theorem applies when z follows a random walk, as shown in Heathcote, Storesletten, and Violante (2017). We deviate from the no-trade conditions by including transfers, which is why we assume a simpler static model. The no-saving assumption is relaxed in the quantitative model of Section 3. In a previous version, we derived an analytical relation between optimal transfers and progressivity with insurable and uninsurable shocks, as in Heathcote, Storesletten, and Violante (2017). While the expression gets cumbersome, the negative relationship between progressivity and transfers remains.

$Y_0(\tau) = n_0(\tau)$, and the tax function level parameter $\lambda_0(\tau)$ from the government budget constraint. Then, for each progressivity τ, we can derive aggregate welfare as the sum of households' utility:

$$W(\tau) = \underbrace{\underbrace{\log\left(n_0(\tau) - G\right)}_{\text{Size}} \underbrace{-\frac{1-\tau}{1+\varphi}}_{\text{Labor disutility}}}_{\text{Efficiency}} \underbrace{-(1-\tau)^2 \frac{v_\omega}{2}}_{\text{Redistribution}}. \tag{6}$$

This expression has a straightforward economic interpretation. The first two terms capture the *efficiency* concerns governing the optimal choice of τ. Because progressivity depresses labor supply, a larger τ reduces the *size* of the economy, and thus aggregate consumption $C_0(\tau) = n_0(\tau) - G$. Yet a larger τ also reduces *labor disutility* $-B\frac{n^{1+\varphi}}{1+\varphi}$, which, using equation (5), equals the second term in the formula.[7] The last term is proportional to the variance of log consumption and captures *redistribution* concerns. Larger progressivity reduces dispersion in consumption, which is welfare improving. As such, heterogeneity increases optimal progressivity.

We use the welfare formula in (6) to numerically illustrate the forces determining optimal progressivity τ. The Frisch elasticity φ^{-1} is set to 0.4, and the labor disutility parameter B is chosen such that labor supply $n_0(\tau) = 0.3$. Progressivity is fixed at $\tau = 0.18$ and government spending is set to $g \equiv G/Y = 23\%$ of output, based on the U.S. estimates of Heathcote, Storesletten, and Violante (2017). We set $v_\omega = 0.268$ to match a variance of log consumption of $v_c = 0.18$, as also measured in Heathcote, Storesletten, and Violante (2017).

Optimal progressivity is zero with a representative agent ($v_\omega = 0$) and no public spending. Positive spending decreases progressivity to $\tau = -0.26$. This negative relation between G and τ relates to the notion of "fiscal pressure" discussed by Heathcote, Storesletten, and Violante (2017) and further developed by Heathcote and Tsujiyama (2021): a planner finances larger spending with less progressive taxes to incentivize labor supply.[8] Finally, adding heterogeneity increases progressivity to $\tau = 0.24$ because of redistribution concerns.

As τ is positive, average and marginal tax rates both increase with income. The presence of transfers will loosen the tight link between average and marginal rates, as we discuss next.

[7] When $v_\omega = 0$, the progressivity τ that maximizes (6) implements the representative-agent first-best allocation. See Appendix A.2 for a derivation of the representative-agent case.

[8] See Ayaz, Fricke, Fuest, and Sachs (2022) for a related discussion.

2.3 Optimal income-tax progressivity with transfers

With non-zero transfers, the policy function for labor does not admit a closed-form solution. Instead, we use the implicit function theorem and approximate the policy around the case with $T = 0$:

$$\hat{n}_i \approx n_0(\tau) - \frac{T}{1+\varphi} \frac{n_0(\tau)}{n_0(\tau) - G} \exp\left(-\tau(1-\tau)\frac{v_\omega}{2}\right) z_i^{-(1-\tau)}. \tag{7}$$

Note that labor supply falls in transfers because of the wealth effect.

Using this approximation, we follow similar steps as in the case without transfers to obtain an expression for welfare as a function of progressivity τ and transfer T:

$$W(\tau, T) = W(\tau, 0) + T \left[\Omega_e(\tau, v_\omega) + \Omega_r(\tau, v_\omega) \right]. \tag{8}$$

The square brackets capture the marginal effect of transfers T on welfare for a given progressivity level τ. As we show below, the marginal effect of transfers decreases with τ, which explains an optimally negative relation between τ and T. This effect can be decomposed in two terms: an *efficiency* term $\Omega_e(\tau, v_\omega)$ and a *redistribution* term $\Omega_r(\tau, v_\omega)$. We discuss each term next.

When $v_\omega = 0$, the term $\Omega_e(\cdot)$ captures an *efficiency* trade-off in the representative-agent economy:

$$\Omega_e(\tau, 0) \equiv \underbrace{u_c(C_0(\tau)) \left.\frac{\partial Y^{ra}(\tau, T)}{\partial T}\right|_{T=0}}_{\text{Size}<0} \underbrace{+u_n(n_0(\tau)) \left.\frac{\partial n^{ra}(\tau, T)}{\partial T}\right|_{T=0}}_{\text{Labor disutility}>0}, \tag{8.a}$$

where Y^{ra} and n^{ra} denote output and labor in the representative-agent economy. Output decreases with transfers, and the welfare cost of smaller output is evaluated using the marginal utility of consumption (*size* effect). At the same time, lower hours induce a welfare gain, evaluated using the marginal utility of leisure (*labor disutility* effect). For $v_\omega = 0$, the term $\Omega_e(\cdot)$ depends on progressivity as follows.

Claim 1 *When $v_\omega = 0$, the welfare gains of transfers associated with efficiency decrease with τ. That is, $\partial \Omega_e(\tau, 0)/\partial \tau < 0$.*

Larger progressivity reduces output, making it costlier to reduce output further with transfers—that is, the *size* cost of transfers becomes larger. In addition, larger progressivity increases leisure, making it less valuable to increase leisure further with transfers—the *disutility* gain of transfers becomes lower. Overall, both progressivity and transfers depress labor supply, making them substitutes in efficiency terms.

When $v_\omega > 0$, the overall *efficiency* term $\Omega_e(\cdot)$ is given as

$$\Omega_e(\tau, v_\omega) = \Omega_e(\tau, 0) + u_c(C_0(\tau)) \left[\left. \frac{\partial Y(\tau, T)}{\partial T} \right|_{T=0} - \left. \frac{\partial Y^{ra}(\tau, T)}{\partial T} \right|_{T=0} \right], \tag{8.b}$$

where the additional term captures the heterogeneity of wealth effects on labor supply. While non-monotonic in τ, this second term is typically small in our calibrations, and $\partial[\Omega_e(\tau, v_\omega)]/\partial\tau < 0$ holds quantitatively so that the total efficiency gains of transfers decrease with progressivity.[9]

Finally, the last term in (8) reflects *redistribution* concerns associated with transfers when agents are heterogeneous:

$$\Omega_r(\tau, v_\omega) = \mathbb{E}\left[u_c(c_0(\tau))\right] - u_c(C_0(\tau)) = \frac{(1-\tau)^2}{n_0(\tau) - G} v_\omega, \tag{8.c}$$

where $c_0(z, \tau)$ is individual consumption, as in equation (2). This term captures dispersion in consumption and is strictly positive as long as $v_\omega > 0$ and $\tau < 1$. Indeed, transfers are welfare improving, as they reduce dispersion in marginal utilities of consumption. The term $\Omega_r(\cdot)$ depends on progressivity as follows.

Claim 2 *Under mild conditions, the welfare gains of transfers associated with redistribution decrease with τ. That is, $\partial\Omega_r(\tau, v_\omega)/\partial\tau \leq 0$.*

This result is intuitive. Larger progressivity already reduces the dispersion in marginal utilities such that there are fewer gains from reducing consumption dispersion even further with higher transfers. Thus, transfers and progressivity act as substitutes to redistribute resources. In the extreme case of $\tau = 1$, after-tax incomes are equalized, and the redistributive gains from having a positive transfer are zero; that is, $\Omega_r(1, v_\omega) = 0$. Overall, *redistribution* concerns strengthen the negative optimal relationship between transfers and tax progressivity driven by *efficiency* concerns: the more generous the transfer T, the smaller the progressivity τ.

We use the welfare formula in equation (8) to evaluate the marginal value of transfers at the illustrative calibration discussed above. We find that *redistribution* gains $\Omega_r = 0.78$ are larger than efficiency concerns $\Omega_e = -0.54$, implying that transfers should be positive in the United States. We further use this formula to compute the optimal progressivity τ for each transfer level T, which we report in Figure 1. As transfers T increase, the optimal income-tax progressivity τ declines.

[9] See Appendix A.2 for a derivation of the representative-agent case and a more detailed discussion of the additional term in (8.b). Appendix A.2 also discusses how the negative relation between transfers and income-tax progressivity derived above compares to the negative relation between spending and income-tax progressivity derived in Heathcote and Tsujiyama (2021).

2.4 Optimal transfers

The welfare formula we derived is insightful on the optimal relationship between T and τ, but it is based on an approximation around $T = 0$. Hence, we now pursue a numerical solution of the model. The goal is twofold: to check the accuracy of the approximation and to compute the optimal level of transfers.

The numerical solution confirms the optimal negative relation between transfers T and progressivity τ around and away from $T = 0$, as Figure 1 shows. In addition, our approximation accurately characterizes the optimal tax progressivity for a wide range of transfers. Figure 1 also plots welfare as a function of transfers. The highest welfare is achieved with a large lump-sum transfer combined with negative progressivity: T amounts to roughly 24% of mean calibrated income, and income taxes are regressive, with $\tau = -0.10$.

Figure 2 plots the average and marginal $t\&T$ rates implied by the welfare maximizing policy. The optimal $t\&T$ system is progressive in terms of average rates because of the optimally large transfers. However, it is regressive in terms of marginal rates because of the optimal negative τ. Thus, the optimal system features higher progressivity in average than in marginal $t\&T$ rates–to achieve redistribution while preserving efficiency.

Comparison to the log-linear plan without a transfer.—Figure 2 also includes the average and marginal $t\&T$ rates of the optimal log-linear plan when transfers are constrained to be zero. In this case, income-tax progressivity is high, at $\tau = 0.24$, because a positive τ is the only tool available for the planner to redistribute. However, in the absence of transfers, redistribution through increasing average rates implies increasing marginal rates. A lump-sum transfer allows to break the tight link between average and marginal rates imposed by the log-linear function, which is welfare improving, as this simple calibration suggests.

Comparison to the Mirrlees plan.—We compute the Mirrlees allocation to derive the second-best nonparametric fiscal plan, which we also plot in Figure 2.[10] The optimal log-linear tax with a transfer comes remarkably close to the Mirrlees plan. Average rates are very similar, and, while marginal rates are non-monotonic under the Mirrlees allocation, marginal rates are comparable over a large domain of the income distribution. Consistently, the optimal log-linear tax with a transfer generates a welfare gain of +0.90% in consumption equivalent terms, almost as large as with the Mirrlees plan (+0.93%). In contrast, the optimal log-linear plan without a transfer generates more modest gains, at only +0.15%.

We further compare the welfare gains delivered under various alternative calibrations. In particular, we present calibrations with a higher Frisch elasticity (as in Heathcote, Storesletten, and Violante 2017), a higher spending-to-output ratio, and with a

[10]See Appendix A.3 for more details on the derivation of the Mirrlees allocation.

Pareto tail for productivity overlayed on the normal shock (as in Heathcote and Tsujiyama 2021).[11] We compute the optimal plan under different Pareto tails. For each level of the Pareto tail, we adjust the variance of the normal shock to match the same overall variance of log consumption, v_c. We report the Pareto coefficient on the consumption distribution, λ_c, and the Pareto coefficient of the income distribution, denoted λ_y, is given by $\lambda_y = (1 - \tau)\lambda_c$. We explore parameterizations for λ_c going from 2.69, as in Heathcote and Tsujiyama (2021), to 3.38, as estimated in Toda and Walsh (2015) using Consumer Expenditure Survey data from 1979 to 2004. Table 1 presents welfare, in consumption-equivalent terms, for: the Mirrlees allocation, the log-linear tax with a transfer, the log-linear tax without a transfer, and an affine tax—that is, a flat income tax with a transfer.

The optimal log-linear tax with a transfer consistently generates welfare gains almost as large as the Mirrlees allocation, regardless of the calibration. This result is intuitive. In this simple model, the government cares about redistributing at the bottom while preserving efficiency at the top. The transfer precisely achieves redistribution at the bottom, while income-tax progressivity controls the slope of marginal rates at the top.

Whether the affine or the log-linear tax with no transfers performs well depends on the calibration of the productivity shock, and in particular, on the thickness of the Pareto tail. As the Pareto tail thickens, the Mirrlees plan features more increasing (i.e. steeper) marginal rates at the top (Mankiw, Weinzierl, and Yagan 2009). This property of the plan can be achieved with a progressive log-linear tax. Instead, when the Pareto tail is thinner, marginal rates are flatter, and the planner uses a larger transfer to provide redistribution. The Mirrlees plan resembles more an affine plan. The log-linear tax with a transfer combines both features, achieving redistribution with transfers while, if needed, implementing increasing marginal rates at the top.

Overall, the log-linear tax has been extensively used in quantitative work because of its tractability, but it imposes significant constraints on the optimal plan. Adding a transfer/intercept to the log-linear tax function retains tractability while consistently generating welfare gains close to the second-best allocation.

2.5 Estimating transfers and income-tax progressivity

While the log-linear tax function fits well the overall $t\&T$ system in the United States (Heathcote, Storesletten, and Violante 2017), we show that adding an intercept improves the empirical fit substantially. Importantly, the fit improves for bottom and top income earners, precisely where it matters most for redistribution and efficiency purposes.

[11] Our setup is simpler that in Heathcote, Storesletten, and Violante (2017), but actually comparable to that in Heathcote and Tsujiyama (2021). See Appendix A.3 for more details.

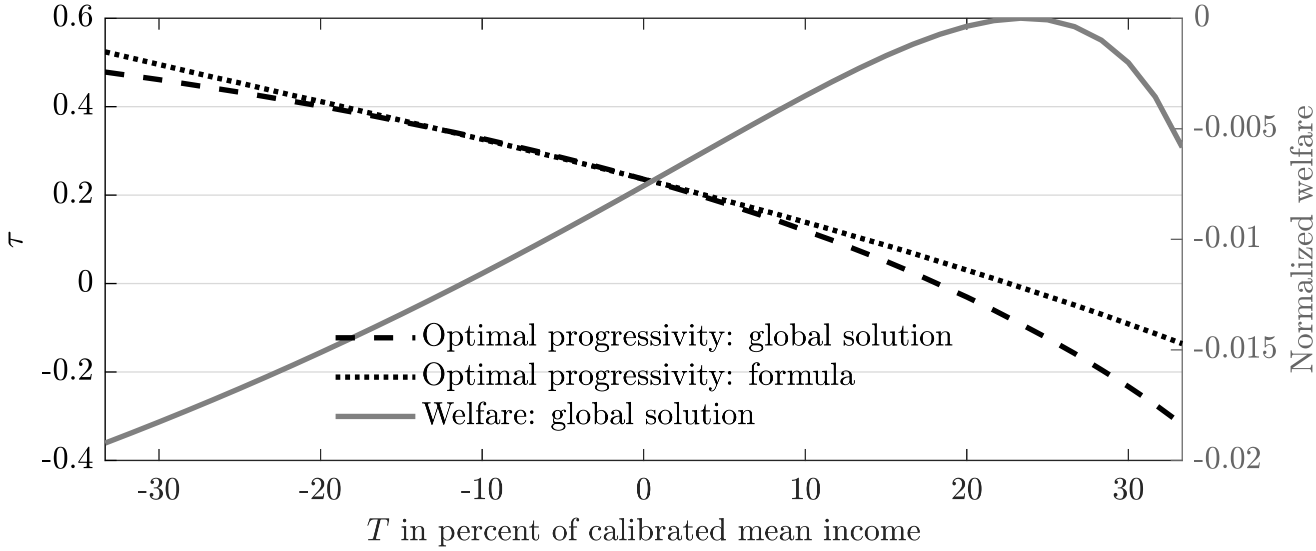

Figure 1: Optimal income-tax progressivity given transfers

Notes: This figure compares, for each transfer, the optimal income-tax progressivity implied by the formula to the global (numerical) solution. It also shows, for each transfer, the welfare associated with the optimal progressivity. We normalize welfare to be zero at the optimal policy.

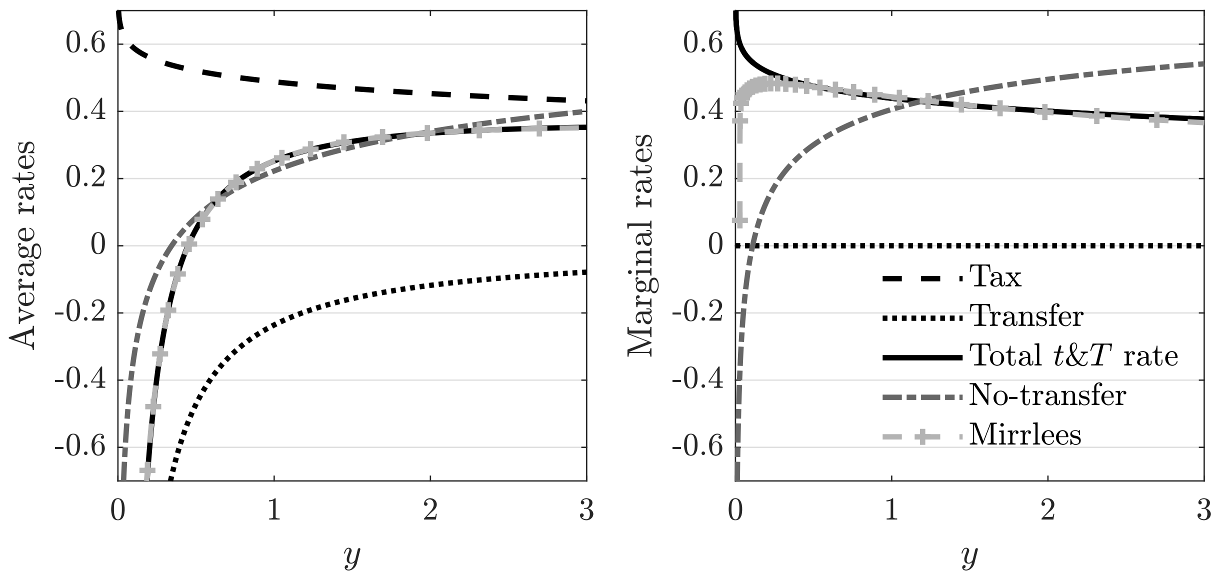

Figure 2: Optimal tax-and-transfer system: Tax rates

Notes: This figure shows average rates (left panel) and marginal rates (right panel) of the optimal tax-and-transfer system. The figure also plots optimal average and marginal rates in: (1) the optimal log-linear plan without a transfer, denoted "No-transfer"; and (2) the Mirrlees allocation, denoted "Mirrlees". Income is normalized by calibrated mean income.

Table 1: Welfare gains: Mirrlees and Ramsey plans

	Mirrlees	Log-linear with T	Log-linear	Affine
Benchmark				
$v_c = 0.18,\ \lambda_c = \infty,\ \varphi = 2.5,\ g = 0.23$	0.93%	0.90%	0.14%	0.84%
Heathcote, Storesletten, and Violante (2017)				
$v_c = 0.18,\ \lambda_c = \infty,\ \varphi = 2.0,\ g = 0.23$	0.81%	0.79%	0.03%	0.70%
Heathcote and Tsujiyama (2021)				
$v_c = 0.23,\ \lambda_c = 2.69,\ \varphi = 2.0,\ g = 0.19$	2.07%	1.97%	1.65%	1.36%
Thinner Pareto tail: $\lambda_c = 2.86$	1.92%	1.85%	1.45%	1.45%
Thinner Pareto tail: $\lambda_c = 3.38$	1.78%	1.74%	1.09%	1.64%

Notes: This table reports, across several calibrations, welfare gains, in consumption-equivalent terms, of implementing: (1) the Mirrlees plan, (2) the optimal log-linear plan with a transfer, (3) the optimal log-linear plan without a transfer, and (4) and the optimal affine plan. v_c denotes the variance of log consumption, λ_c the thickness of the Pareto tail of the distribution of consumption, and g the spending-to-output ratio, in the respective calibrations.

In particular, we present three estimates. First, we estimate the log-linear tax function with no transfers, as in Heathcote, Storesletten, and Violante (2017). This function can be easily estimated using ordinary least squares by regressing the log of after-$t\&T$ income on the log of pre-$t\&T$ income. Second, we estimate the same tax function but use income in levels and thus nonlinear least squares. Third, we add a transfer, also using income in levels and nonlinear least squares. We construct pre-$t\&T$ and after-$t\&T$ income at the household level using Current Population Survey (CPS) data for 2013. Our measure of taxes and transfers includes: personal income federal and state taxes; payroll taxes; tax credits, the most important being the Earned Income Tax Credit (EITC) and the Child Tax Credit (CTC); Supplemental Nutrition Assistance Program (SNAP) and housing assistance, imputed following the Congressional Budget Office (CBO) procedure; and additional welfare transfers, as reported in the CPS. More details can be found in Appendix B.1. For each estimate, the left panel in Figure 3 reports the difference between actual and predicted $t\&T$ paid. The right panel reports the same difference relative to pre-$t\&T$ income. Results are presented by 2.5% bins of pre-$t\&T$ income distribution.

The first estimate delivers a progressivity of $\tau = 0.18$—as in Figure I.A in Heathcote, Storesletten, and Violante (2017)—and fits the U.S. $t\&T$ function remarkably well, especially for a two-parameter tax function. Yet, it overestimates taxes paid at the top and underestimates transfers at the bottom. Prediction errors in taxes paid for the top income decile range from \$10,000 to \$50,000, while errors in transfers received at the bottom 5% range from 20% to 40% of their pre-$t\&T$ income. The errors at the top may have significant implications for the government's revenues, while differences at the

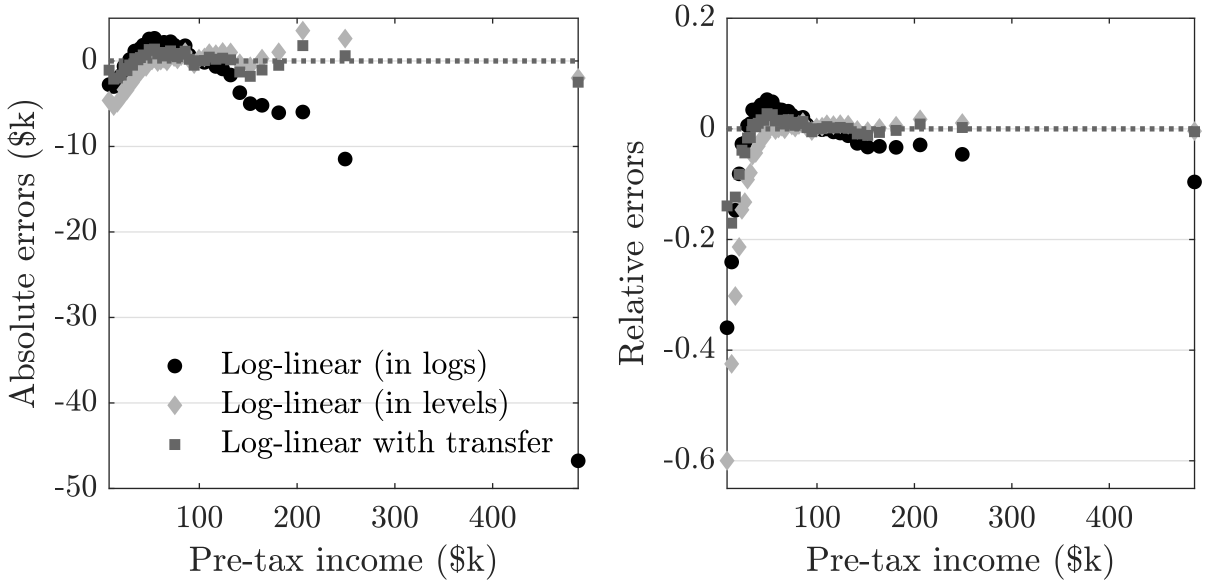

Figure 3: Estimated $t\&T$ system: Prediction errors

Notes: Prediction errors in taxes paid minus transfers received, data minus prediction. Errors are reported for three estimates: (1) the log-linear income-tax function without a transfer, estimated on log income; (2) the log-linear income-tax function without a transfer, estimated on income in levels; and (3) the log-linear income-tax function with a transfer, estimated on income in levels. The left panel reports absolute errors in dollars, while the right panel reports relative errors as a fraction of pre-$t\&T$ income, by 2.5% bins of pre-$t\&T$ income distribution.

bottom may amplify redistribution concerns in the status quo. Estimating the log-linear tax function in levels improves the fit at the top, with estimated progressivity falling to $\tau = 0.09$, but it deteriorates the fit at the bottom. Adding an intercept allows to match both taxes paid at the top and transfers received at the bottom. The estimated parameters are $\tau = 0.06$ and $T = \$4{,}500$.

2.6 Taking stock

There are two main takeaways from this section. First, adding a transfer to the log-linear tax function results in a substantial increase in welfare, which comes consistently close to the Mirrlees plan. This is because a transfer allows to implement more progressive average than marginal rates. Second, adding a transfer also significantly improves the empirical fit of the log-linear tax function on U.S. data.

While calibrated, the model of this section is too simple to make a truly quantitative statement about the optimal combination of the instruments. Furthermore, a lump-sum transfer is a limited description of currently available instruments. To address these points, we move next to a dynamic quantitative macro model of the U.S. economy with more flexible fiscal instruments.

3 A quantitative model

We extend our previous discussion by analyzing the optimal design of taxes and transfers in a rich dynamic quantitative environment with endogenous self-insurance. In particular, we incorporate a flexible $t\&T$ system into a canonical heterogeneous-agent model (Aiyagari 1994) augmented with realistic labor income risk. We describe the model environment and its calibration in this section, including the empirical fit of our flexible tax function on U.S. household-level data. Section 4 discusses the optimal tax plan and the role of transfers.

3.1 Environment

The economy is populated by a continuum of households, a representative firm, and a government. The firm produces output by combining labor and capital, both of which are supplied by households. The government finances transfers and spending by taxing households' labor and capital incomes as well as consumption. We present the economy at its stationary equilibrium but will consider transitions when evaluating tax reforms.

Households.—Households value consumption c and leisure $1-n$. Their idiosyncratic labor productivity z follows a Markov process with transition probabilities $\pi_z(z', z)$. Labor productivity shocks are uninsurable: households can only trade a one-period risk-free bond to self-insure, subject to a borrowing limit. Let $V(a, z)$ be the maximal attainable value to a household with assets a and idiosyncratic productivity z:

$$\begin{aligned} V(a,z) = \max_{c,a',n} & \left\{ \frac{c^{1-\sigma}}{1-\sigma} - B\frac{n^{1+\varphi}}{1+\varphi} + \beta \mathbb{E}_{z'}\left[V\left(a',z'\right)|z\right] \right\} \\ & \text{s.t.} \\ & (1+\tau_c)c + a' \leq wzn + (1+r)a - \mathcal{T}\left(wzn, ra\right) \\ & a' \geq \underline{a}, \end{aligned} \tag{10}$$

where w and r stand for wages and the interest rate, respectively; $\underline{a}$ denotes the borrowing constraint; and τ_c is a flat consumption tax. Households' income taxes and transfers are captured by $\mathcal{T}(wzn, ra)$, which depend on labor income wzn and capital earnings ra. We discuss the shape of $\mathcal{T}(\cdot)$ in detail below. Let $n(a, z), c(a, z)$ and $a'(a, z)$ denote a household's optimal policies.

Representative firm.—The representative firm demands labor and capital in order to maximize current profits

$$\Pi = \max_{K,L} \left\{ K^{1-\alpha} L^{\alpha} - wL - (r+\delta)K \right\}, \tag{11}$$

where δ is the depreciation rate of capital. Optimality conditions for the firm are standard:

marginal products are equalized to the cost of each factor.

Government.—The government's budget constraint is given by

$$G + (1+r)D = D + \int \mathcal{T}(wzn, ra)d\mu(a,z) + \tau_c \int c(a,z)d\mu(a,z) \tag{12}$$

where G is government spending, D is government debt, and $\mu(a,z)$ is the measure of households with state (a,z) in the economy.

Stationary equilibrium.—Let A be the space for assets and Z the space for productivity. Define the state space $S = A \times Z$, and let $\mathcal{B}$ be the Borel $\sigma-$algebra induced by S. A formal definition of the competitive stationary equilibrium for this economy is provided next.

A **competitive stationary equilibrium** for this economy is given by value function $V(a,z)$ and policies $\{n(a,z), c(a,z), a'(a,z)\}$ for the household; policies for the firm $\{L, K\}$; government decisions $\{G, D, \mathcal{T}, \tau_c\}$; a measure μ over $\mathcal{B}$; and prices $\{r, w\}$ such that, given prices and government decisions: *(i)* households' policies solve their problems and achieve value $V(a,z)$, *(ii)* the firm's policies solve its problem, *(iii)* the government's budget constraint is satisfied, *(iv)* the capital market clears: $K + D = \int_{\mathcal{B}} a'(a,z)d\mu(a,z)$, *(v)* the labor market clears: $L = \int_{\mathcal{B}} zn(a,z)d\mu(a,z)$, *(vi)* the goods market clears: $Y = \int_{\mathcal{B}} c(a,z)d\mu(a,z) + \delta K + G$, and *(vii)* the measure μ is consistent with households' policies: $\mu(\mathcal{B}) = \int_{\mathcal{B}} Q((a,z),\mathcal{B})d\mu(a,z)$, where Q is a transition function between any two periods defined by $Q((a,z),\mathcal{B}) = \mathbb{I}_{\{a'(a,z)\in\mathcal{B}\}} \sum_{z'\in\mathcal{B}} \pi_z(z',z)$.

3.2 A flexible tax-and-transfer function

We endow the government with two fiscal tools capturing the key elements of the U.S. *t&T* system: a nonlinear labor tax and targeted transfers. Conveniently, the overall *t&T* system is characterized by a few parameters only, all of which have a clear economic intuition. Yet the function allows for flexible—potentially non-monotonic—shapes of the overall marginal *t&T* rates, a feature often found desirable in the optimal taxation literature.

In particular, we divide the *t&T* function $\mathcal{T}(\cdot)$ into three components: a flat tax τ_k on capital income y_k, a nonlinear tax $\tau(y_\ell)$ on labor income y_ℓ, and a targeted transfer component $T(y)$ on total income $y = y_k + y_\ell$. From these components, we keep the capital tax constant (and purposely simple) and focus attention on labor taxes and transfers.

We assume the average labor tax rate is characterized by two parameters, θ and λ, as

$$\tau(y_\ell) = \exp\left(\log(\lambda)\left(\frac{y_\ell}{\bar{y}}\right)^{-2\theta}\right), \tag{13}$$

where $\bar{y}$ is mean income. As with the log-linear tax function used in Section 2, our

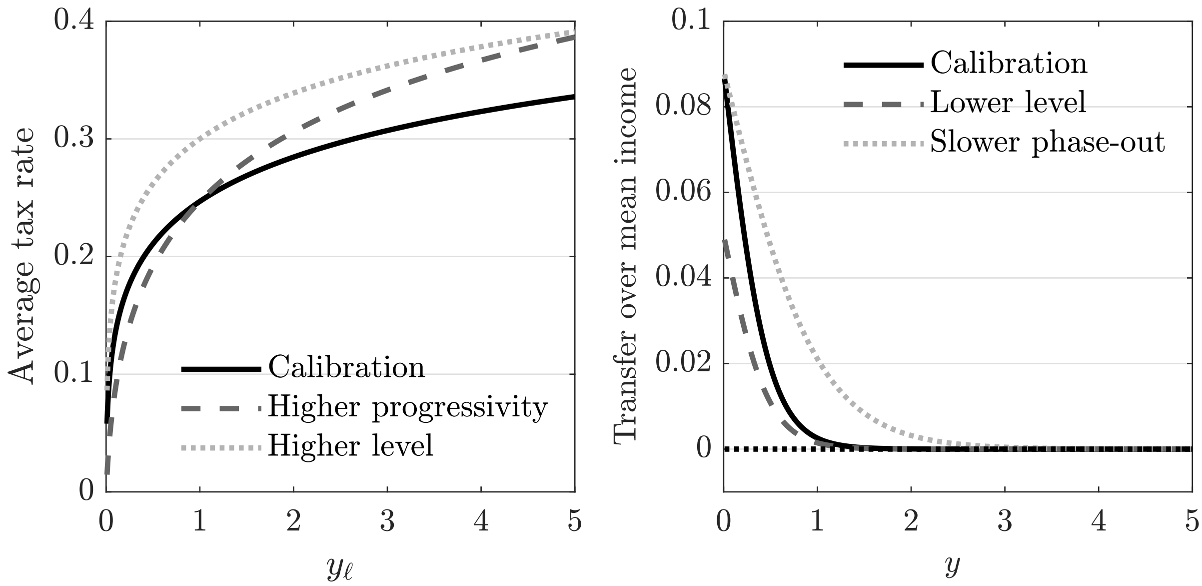

Figure 4: Labor tax function and targeted transfers

Notes: The left panel illustrates the shape of the new labor tax function. It compares the average tax rate in the calibration ($\lambda = 0.25$ and $\theta = 0.08$) to a higher progressivity ($\theta = 0.12$) and a higher level ($\lambda = 0.3$). The right panel plots the new transfer function. It compares transfers normalized by mean income in the calibration ($m = 0.09$ and $\xi = 4.22$) to a lower level ($m = 0.05$) and a slower phase-out ($\xi = 2$). Income is normalized by mean income.

proposed tax function has two interpretable parameters: θ for the progressivity and λ for its level. A positive (negative) θ implies marginal tax rates that increase (decrease) with income. At $\theta = 0$, the tax is flat at λ. For all θ, the tax rate is exactly λ when income is at its mean, $y_\ell = \bar{y}$. The left panel of Figure 4 shows how $\tau(y_\ell)$ varies with θ and λ.

This labor tax function is always non-negative, in line with statutory tax rates in the United States. Other than that, our function largely resembles the log-linear tax used in Section 2, as shown in Figure C.1 of Appendix C.1. The level of progressivity θ is approximately on the same scale as the progressivity τ of the log-linear tax. The non-negative labor taxes imply that we rely exclusively on income-dependent transfers to generate negative $t\&T$ rates, as we explain next.

We assume a transfer function that is characterized by two parameters: a level m and a phase-out rate ξ. In particular, the transfer given to a household with total income y is given as

$$T(y) = m\bar{y}\frac{2\exp\left\{-\xi\left(\frac{y}{\bar{y}}\right)\right\}}{1+\exp\left\{-\xi\left(\frac{y}{\bar{y}}\right)\right\}}. \tag{14}$$

The parameter m measures transfers to a household with zero income as a multiple of mean income $\bar{y}$. The parameter ξ determines how quickly transfers phase out with total income. When $\xi = 0$, transfers are a lump sum. As ξ becomes larger, transfers phase out faster. The right panel of Figure 4 shows how transfers vary with m and ξ.

This functional form for transfers is a realistic description of U.S. income security programs, where transfers are means-tested, typically on both labor and capital income. Moreover, the phasing out of the transfer allows for non-monotonic marginal rates of the entire $t\&T$ system, which is not possible with a lump-sum transfer.

Thus, a government's policy is characterized by four parameters: the progressivity and level of labor taxes, θ and λ, and the level and phase-out rate of transfers, m and ξ. We refer to a fiscal plan $\boldsymbol{\tau} = \{\theta, \lambda, m, \xi\}$ as a government's policy that satisfies its budget constraint (12).

3.3 Estimation of the tax-and-transfer function

We estimate the tax and the transfer functions using CPS household-level data, independently of other model parameters. Our measure of taxes includes personal federal and state income taxes, as well as employer and employee payroll taxes, from which we deduct the non-refunded part of federal and state tax credits. Our measure of transfers includes refunded tax credits—that is, the (refunded) Additional Child Tax Credit (ACTC) and the refunded part of the EITC and of state credits—as well as the SNAP, Housing Assistance, and additional welfare transfers. Appendix B.1 contains details on the construction of variables and robustness checks of the estimates, as well as the distribution of all components of transfers across households (Figure B.1).

Figure 5 plots estimated taxes and transfers in dollars against their data counterpart. Income-tax progressivity is estimated at $\theta = 0.08$, close to the estimate obtained in Section 2.5 with a log-linear function and a lump sum. Transfers for a household at minimum income are estimated at \$7,100, with a phase-out of $\xi = 4.22$. Overall, our flexible functions require only a small number of parameters and fit remarkably well the distributions of taxes and transfers, making it particularly useful for applied/quantitative work.[12]

3.4 Calibration

We calibrate the model to the U.S. economy in 2012. We take a period in the model to be a year. Several parameters are calibrated within the model, while others are taken from the literature, as we discuss next.

Labor income risk.—Recent empirical work has unveiled key statistics about labor income dynamics, two of which are potentially important for our purposes. First, the recent work in Guvenen, Karahan, Ozkan, and Song (2021) shows that earnings growth rates are negatively skewed and exhibit excess kurtosis. That is, relative to a normal distribution, there are more individuals with small and large earnings changes, but fewer

[12] We estimate labor taxes and transfers separately. Figure B.2 in Appendix B.1 shows that prediction errors of the after-$t\&T$ income are systematically small across the income distribution.

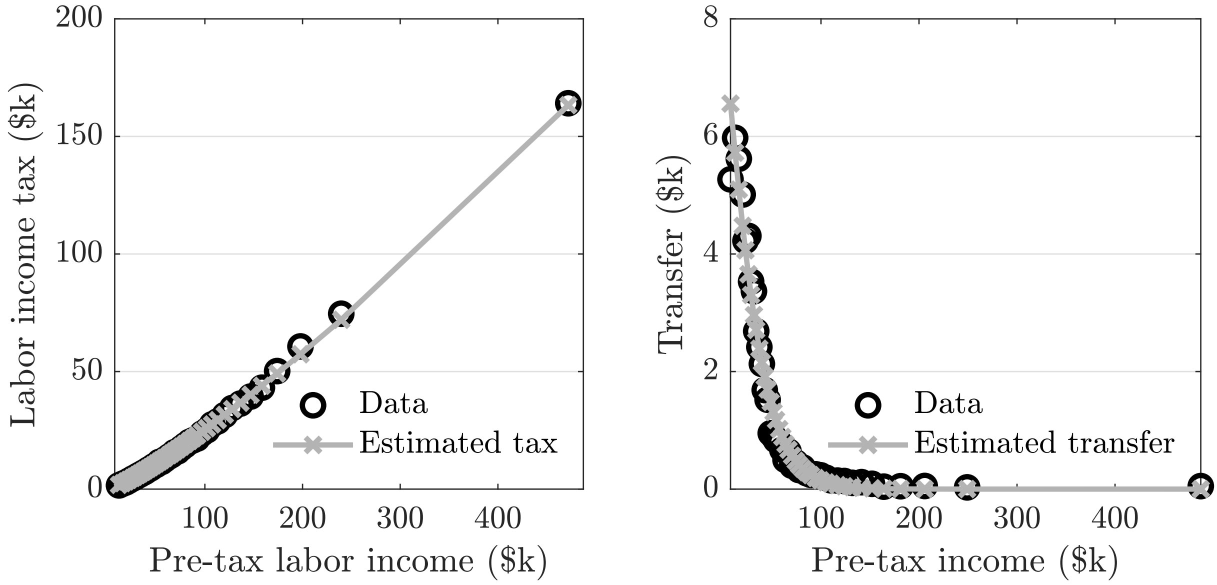

Figure 5: Estimated tax and transfer functions

Notes: The left panel plots taxes paid, while the right panel plots transfers received, by 2.5% bins of pre-$t\&T$ income distribution.

with medium-sized earnings changes. Often, the large earnings changes are negative. Second, labor income inequality has increased in recent decades. These facts are relevant to the efficiency and redistribution tradeoffs we analyze in this paper. Thus, we propose a process for idiosyncratic labor risk that can account for them.

In particular, we assume that a household's productivity z follows a Gaussian Mixture Autoregressive (GMAR) process in logs as

$$\begin{aligned} \log z_t &= \rho \log z_{t-1} + \eta_t \\ \eta_t &\sim \begin{cases} \mathcal{N}\left(\mu_1, \sigma_1^2\right) & \text{with probability } p_1, \\ \mathcal{N}\left(\mu_2, \sigma_2^2\right) & \text{with probability } 1 - p_1, \end{cases} \end{aligned} \tag{15}$$

with $\mathbb{E}[\eta] = p_1\mu_1 + (1 - p_1)\mu_2 = 0$, so that μ_2 is pinned down given μ_1 and p_1.

Most times, households draw an innovation from the first normal that has a low variance. Draws from the second normal occur infrequently, but they have a large variance and a negative mean. In this way, we can generate frequent small and infrequent large negative earnings changes—that is, the negative skewness and excess kurtosis empirically documented for labor income growth. We discretize the productivity process using the method of Farmer and Toda (2017).

Furthermore, to better capture the concentration of incomes at the top, we follow Hubmer, Krusell, and Smith (2020) and make one additional adjustment to the productivity process. We adjust the top 15% states in our productivity grid such that they follow a Pareto distribution with tail $\kappa = 1.6$, as estimated by Aoki and Nirei (2017).

The income process is then characterized by five parameters: $(\rho, \mu_1, \sigma_1, \sigma_2, p_1)$. We pick these parameters to match key statistics of households' labor earnings growth from the Panel Study of Income Dynamics (2021) (PSID) as well as income concentration at the top.[13] We target four moments of the labor income growth distribution: the standard deviation (0.33), the difference between the ninetieth and the tenth percentiles (0.60), the skewness (-0.37), and the kurtosis (12.48). Additionally, we target a labor income share of 41% for the top 10% of labor income earners, as found in the Survey of Consumer Finances (SCF) for 2012.[14]

Government.—We use data from National Income and Product Accounts (NIPA) and calibrate capital taxes to match capital tax revenues over GDP. Similarly, we calibrate consumption taxes to match revenues from sales and excise taxes over total consumption expenditure.[15] This procedure yields $\tau_k = 29.8\%$ and $\tau_c = 6.3\%$. We calibrate government debt D to match a debt-to-output ratio of 99%, as in the Flow of Funds.[16] The parameters for labor taxes and targeted transfers, $\boldsymbol{\tau} = \{\theta, \lambda, m, \xi\}$, are estimated using cross-sectional data, as explained in Section 3.3. Finally, spending is implied by government budget clearing. Overall, total tax revenues amount to 23.6% of GDP, a number close to its data counterpart (NIPA). Tax revenues are split as follows: spending sums up to 20.4% of GDP, transfers to 1.2% of GDP, and interest payments to 2.0% of GDP.[17] Table 3 reports the distribution of transfers and labor income-tax rates across households, while Figure 6 plots average and marginal $t\&T$ rates as a function of labor income for different values of capital income.

Remaining parameters.—We set the coefficient of relative risk aversion σ to 2, a value that is more standard in quantitative macroeconomics than the log utility used for tractability in Section 2. We fix the Frisch elasticity φ to 0.4, also a common value. We set the production side parameters to standard values, with the labor share $\alpha = 0.64$ and the (annual) depreciation rate $\delta = 0.06$. The borrowing constraint is calibrated to a quarter of mean income. We calibrate the discount factor β and labor disutility B to jointly match an interest rate of 2% and an average labor supply of 0.3. All parameters are summarized in Table 2. Section 5 presents robustness with respect to several parameters.

Distributions.—The model matches well the income distribution, as Table 4 shows. While the labor income share of the top 10% is targeted, the model features a remarkable

[13]We do not directly use the statistics reported in Guvenen, Karahan, Ozkan, and Song (2021), because they are computed at the individual level. Instead, we follow De Nardi, Fella, and Paz-Pardo (2019) and use PSID data to compute income growth statistics at the household level, but base estimates on pre-tax income data. See Appendix B.3 for more details on the PSID. Table C.1 in Appendix C.3 reports income risk moments in the model.

[14]See Appendix B.2 for more details on the SCF.

[15]See Appendix B.4 for definitions of capital tax revenues and consumption taxes in the NIPA tables.

[16]See Appendix B.5 for more details on the Flow of Funds data.

[17]We find a comparable split using NIPA data. See Appendix B.4.

Table 2: Parameter values

Household			**Income Process**		
β	Discount factor	0.968	ρ	Persistence	0.935
σ	Risk aversion	2.000	p_1	Weight on first normal	0.850
$1/\varphi$	Labor supply elasticity	0.400	μ_1	Mean of first normal	0.017
B	Disutility of labor	85.00	μ_2	Mean of second normal	-0.096
$\underline{a}$	Borrowing constraint	-0.220	σ_1	Std. dev. of first normal	0.166
Production			σ_2	Std. dev. of second normal	0.535
δ	Depreciation rate	0.060	κ	Pareto tail parameter	1.600
α	Labor share	0.640			
Government					
θ	Tax progressivity	0.077	D	Public debt	1.064
λ	Tax level	0.247	G	Government spending	0.218
m	Transfer level	0.088	τ_k	Capital tax rate	0.298
ξ	Transfer phase-out	4.220	τ_c	Consumption tax rate	0.063

Table 3: Calibration: Taxes and transfers

Tax rates	Q1	Q2	Q3	Q4	Q5
Data	16%	19%	22%	24%	29%
Model	16%	18%	19%	20%	27%
Transfers rates	Q1	Q2	Q3	Q4	Q5
Data	23%	4%	1%	0%	0%
Model	17%	4%	1%	0%	0%

Notes: Total labor income taxes paid and total transfers received over total income per quintile. Data: CPS 2013, working-age households.

Table 4: Calibration: Income and wealth

Labor income	Q1	Q2	Q3	Q4	Q5	Top 10
Data	4%	8%	13%	19%	56%	41%
Model	5%	9%	14%	20%	52%	38%
Net worth	Q1	Q2	Q3	Q4	Q5	Top 10
Data	-1%	1%	3%	10%	87%	75%
Model	-0%	2%	7%	18%	73%	53%

Notes: Labor income shares by labor income quintiles and wealth shares by wealth quintiles. Data: SCF 2013, working-age households.

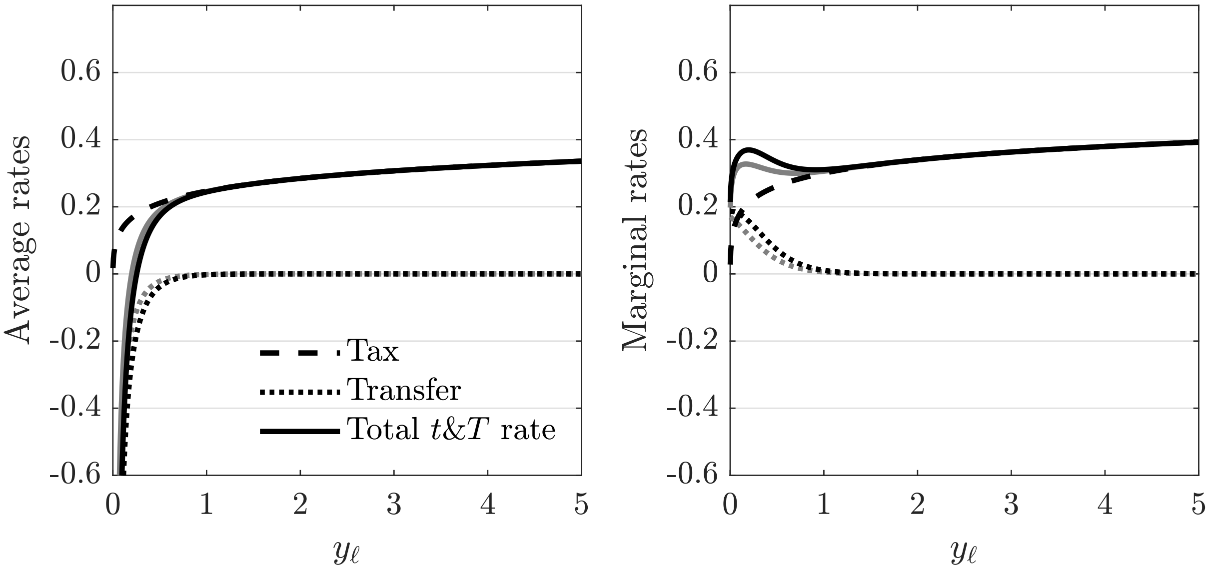

Figure 6: Calibration: Average and marginal tax-and-transfer rates

Notes: Average and marginal $t\&T$ rates in the calibrated steady state along the labor income distribution, normalized by mean income. As transfers phase out with respect to total income, we report two cases: at zero capital income (legend color) and at mean capital income (lighter color).

fit for all income quintiles. Importantly, the model captures accurately how poor the poor are: labor income at the first, fifth, and tenth percentiles amounts to 12%, 24%, and 36% of median labor income, versus 10%, 21%, and 31% in the CPS. The lowest-income household earns \$5,600 in the model, in line with our data earnings threshold of \$5,000. As we show in Section 5.1, the left and right tails of the income distribution shape the optimal plan. In contrast, the model falls short in generating enough wealth concentration at the top, a common shortcoming of this type of model.

Finally, we investigate the dispersion of hours and consumption across households. The variance of log hours at the household level amounts to 0.097 in the data, versus 0.063 in the model. A part of the discrepancy may be explained by measurement error in reported hours.[18] Regarding consumption, the model generates a variance of log consumption of 0.252, a number well in line with typical estimates in the literature (Heathcote, Perri, and Violante 2010, Attanasio and Pistaferri 2014, Heathcote and Tsujiyama 2021).

3.5 Wealth effects and elasticities

Wealth effects on labor supply are key to the optimal design of the $t\&T$ system. Jointly with the Frisch elasticity, wealth effects determine the response of labor supply to a tax change. Golosov, Graber, Mogstad, and Novgorodsky (2021) estimate these wealth effects by combining U.S. taxpayers' income data with data on lottery winnings, a plausibly

[18] Appendix B.1 presents more details on the construction of the measure of hours in the data, while Appendix C.4 report more moments of the distribution of log hours.

exogenous variation in wealth. We replicate their estimates in our model—that is, we assume an unexpected windfall and compute labor income responses.

The model wealth effects are in line with the estimates in Golosov et al. (2021). We report wealth effects for an average and a large win size, the latter being closer to the magnitude of transfers we analyze in Section 4. The estimates represent the five-year average change in labor income for every $100 won, as reported in Golosov et al. (2021). For the average win size, labor income declines $2.3 in the data and $2.4 in the model. For the large win size, the declines are $0.9 in the data and $1.2 in the model.[19]

Finally, the model also generates a reasonable labor elasticity at the top. We conduct a partial-equilibrium experiment in which marginal tax rates unexpectedly increase by 1% for all households. We vary the persistence of the tax change, from a one-year temporary tax reform to a permanent tax reform. The implied labor elasticity of the top 1% varies from 0.12 to 0.34 in the model, a number well within the range of values reported in the literature.[20]

4 Optimal tax-and-transfer plan

In this section, we show that the insights of the analytical model carry over to our quantitative environment: a planner optimally trades higher transfers for lower income-tax progressivity. When including realistic efficiency and redistribution concerns, the optimal fiscal plan features generous transfers and moderate labor-tax progressivity. Thus, as in the analytical case, transfers are key to generate more progressive average than marginal $t\&T$ rates. Furthermore, we show that while a phasing out of transfers is optimal, a lump-sum transfer comes close in terms of welfare gains.

4.1 A Ramsey approach

A government's plan is fully characterized by $\boldsymbol{\tau} = \{\theta, \lambda, m, \xi\}$, the progressivity and level of labor taxes, and the level and phase-out rate of transfers. We use a utilitarian welfare criterion to evaluate a one-time change in policy $\boldsymbol{\tau}$, keeping capital and consumption taxes constant at their calibrated values, and include transitions in the welfare computations.[21]

[19] See Appendix C.5 for more details.

[20] See Kindermann and Krueger (2022) for a discussion of the empirical literature and the model counterpart in a related framework. In a recent paper, Rauh and Shyu (2022) use administrative microdata from California to estimate income responses from a 2012 tax reform to state marginal income tax rates for upper-income households. Their estimates—which are significantly larger than the rest of the literature (Chetty 2012) and thus larger than our model—may also reflect rent-seeking behaviors or fiscal optimization, as the authors mention.

[21] We actually optimize on three parameters—θ, λ, and ξ—and set m to satisfy the budget constraint. Thus, m varies somewhat along the transition. We report the long-run value of m. See Appendix C.2 for computational details. For optimal time-varying tax systems, see Acikgöz, Hagedorn, Holter, and Wang (2022) and Dyrda and Pedroni (2022).

In particular, let $V_0(a, z; \boldsymbol{\tau})$ be the lifetime utility of a household with assets a and productivity z in the period when the policy $\boldsymbol{\tau}$ is implemented. The utilitarian welfare criterion $\mathcal{W}(\boldsymbol{\tau})$ considers the sum of utilities $V_0(\cdot)$ as

$$\mathcal{W}(\boldsymbol{\tau}) = \int V_0(a, z; \boldsymbol{\tau}) d\mu_0(a, z). \tag{16}$$

Notice that policy $\boldsymbol{\tau}$ affects household lifetime utility, but the measure $\mu_0(a, z)$ is given by the initial steady state of the economy.[22]

4.2 Optimal tax-and-transfer system

The optimal system is substantially more redistributive than the system currently in place in the United States. Optimal transfers are large, at $19,800 in 2012 U.S. dollars for the lowest-income household. This value implies an income floor of 23% of mean income ($m = 0.23$). The optimal phase-out is slow—at $\xi = 3.41$, compared to $\xi = 4.22$ in the calibration—implying a transfer of $3,700 for a household at calibrated median income. Optimal income-tax progressivity is moderate at $\theta = 0.14$, only slightly higher than in the status quo.

Optimal average $t\&T$ rates are more progressive than marginal rates, as Figure 7 shows. Average rates monotonically increase with income because of both transfers and progressive labor taxes. Marginal rates, however, are not monotonic. They are high for low-income earners because transfers phase out, lower for medium-income earners, and high again for top-income earners because of progressive labor taxes.

Transfers in the optimal system are substantially more generous than in the status quo. As shown in Table 5, the transfer rate is 85% for the bottom income quintile, compared to only 23% in the data. Transfers also remain generous for the second and third quintiles in the optimal plan, while they are virtually zero in the empirical counterpart. The larger transfers are financed with higher labor taxes, with tax revenues equal to 33% of GDP in the new steady state, but the optimal income-tax progressivity increases only slightly, as compared to the status quo. Thus, tax rates increase almost uniformly across quintiles. Table 5 also shows average and marginal $t\&T$ rates by income quintile. Average $t\&T$ rates are more progressive than in the status quo: they are equal to -66% for the bottom quintile and monotonically increase with income to reach 36% for the top quintile. Marginal $t\&T$ rates are non-monotonic, at above 60% in the bottom two quintiles and around 50% in the top two quintiles.

Overall, the optimal plan achieves redistribution via large transfers at the bottom. It also preserves efficiency with moderate income-tax progressivity, thereby incentivizing labor from productive households. This quantitative finding is in line with the ana-

[22] Appendix D.1 reports the optimal fiscal plan when optimizing on steady-state welfare.

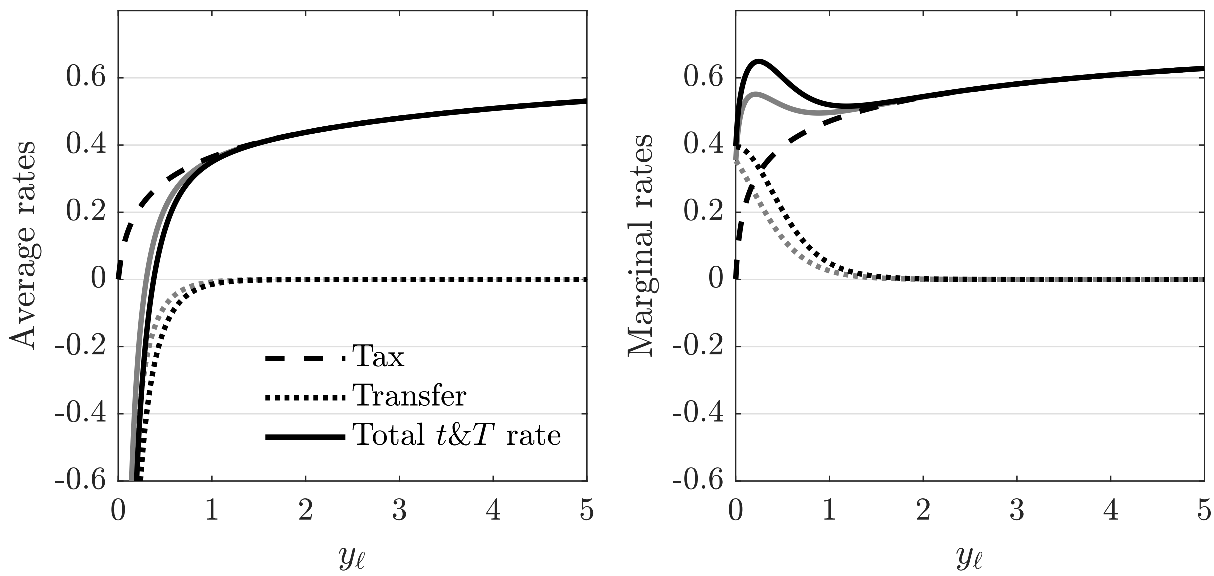

Figure 7: Optimal tax-and-transfer system: Average and marginal $t\&T$ rates

Notes: This figure shows optimal average and marginal $t\&T$ rates along the labor income distribution, normalized by calibrated mean income. We report two cases: zero capital income (legend color) and mean capital income (lighter color).

lytical model, where transfers are optimally used to disentangle average and marginal progressivity.

4.3 The economy shrinks, but most people benefit

The optimal $t\&T$ system results in a decline in economic activity, as Figure 8 shows. The higher taxes and more generous transfers lead to lower savings, and thus wages decline while the interest rate increases. The distribution of hours worked in the final steady state also shifts left, as Table 6 shows. Overall, output in the final steady state is 14% lower than in the initial steady state, a drop that is magnified by the consideration of transitions.[23]

Yet not only does the optimal $t\&T$ system increase utilitarian welfare, but it is also favored by a majority of households over the status quo. Aggregate welfare gains amount to 6.00% in consumption-equivalent terms.[24] As Figure 9 shows, these large welfare gains accrue primarily to the poor, who benefit the most from the generous transfers. Households with high productivity experience welfare losses. On average, though those with higher assets may still benefit from the tax reform because of the higher interest rates. Overall,

[23] When maximizing steady-state welfare, the planner ignores transition gains from price dynamics and lower capital stock, as explained in Bakış, Kaymak, and Poschke (2015). Consequently, the optimal system provides less redistribution, and output falls only by 5%. Appendix D.1 provides a complete description of the fiscal system maximizing steady-state welfare.

[24] See Appendix D.2 for a formal derivation of consumption equivalents in our environment and for more details on the welfare decomposition.

Table 5: Optimal tax-and-transfer system

Data	Q1	Q2	Q3	Q4	Q5
Tax rates	16%	19%	22%	24%	29%
Transfer rates	23%	4%	1%	0%	0%
Average $t\&T$ rates	-7%	16%	21%	24%	29%
Optimal plan	Q1	Q2	Q3	Q4	Q5
Tax rates	19%	23%	26%	26%	36%
Transfer rates	85%	24%	8%	2%	0%
Average $t\&T$ rates	-66%	-1%	19%	24%	36%
Marginal $t\&T$ rates	63%	61%	55%	48%	50%
UBI plan	Q1	Q2	Q3	Q4	Q5
Tax rates	48%	46%	47%	41%	48%
Transfer rates	97%	49%	33%	22%	9%
Average $t\&T$ rates	-48%	-3%	15%	19%	40%
Marginal $t\&T$ rates	56%	58%	59%	59%	61%
Log-linear plan	Q1	Q2	Q3	Q4	Q5
Average $t\&T$ rates	-11%	8%	18%	25%	47%
Marginal $t\&T$ rates	23%	37%	45%	49%	62%

Notes: This table shows per-quintile tax and transfer rates: total labor-income taxes paid and total transfers received over total income, by income quintile. It also reports the average $t\&T$ rates, that is, the tax rate minus the transfer rate; and the marginal $t\&T$ rate, averaged by quintile. It reports the CPS data, the optimal plan with targeted transfers, the UBI plan with lump-sum transfers, and the log-linear plan without transfers.

76% of households benefit from implementing the optimal plan.

We follow Bhandari, Evans, Golosov, and Sargent (2022) and decompose the welfare gains into three components: aggregate efficiency, redistribution, and insurance. Efficiency captures the welfare gains resulting from changes in aggregate resources. Redistribution captures changes in ex-ante shares of consumption and leisure, while insurance captures changes in ex-post utility risk.

About three-fourths of the gains come from insurance, one-fifth from redistribution, and the remainder from the efficiency component. The insurance gains reflect lower volatility of consumption due to larger transfers. All households record a gain in the insurance component, but this gain is larger for low-asset households. The redistribution component is driven by lower dispersion in ex-ante consumption shares. This component aggregates to a small number but masks substantial heterogeneity: it is large for low-asset/low-productivity households, while it is negative for high-productivity but low-asset

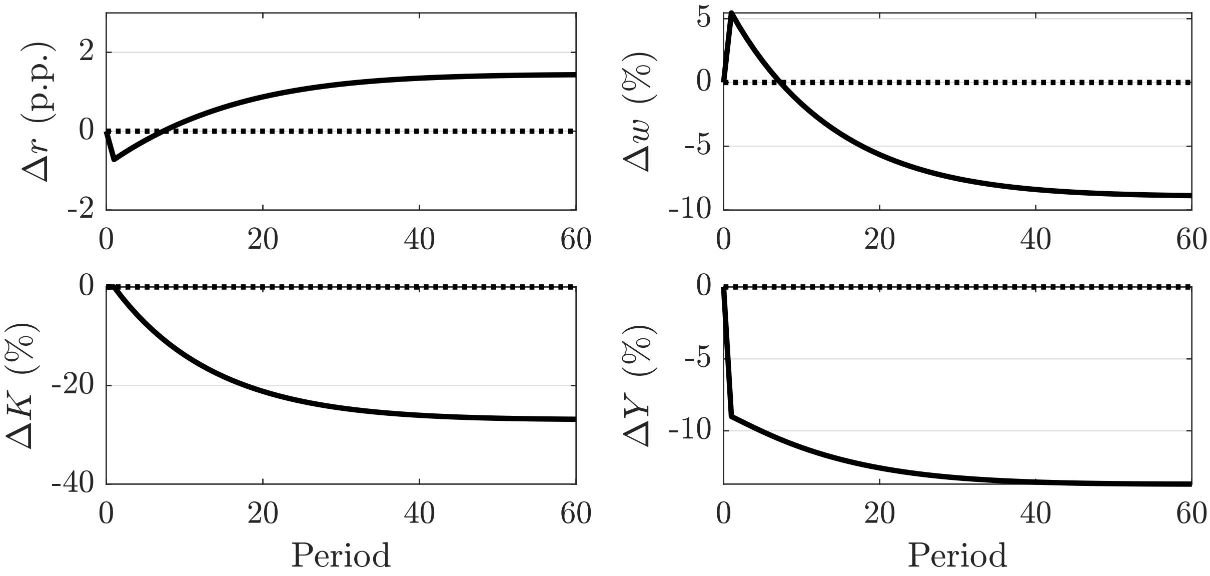

Figure 8: Quantitative model: Responses of prices and quantities to the tax reform

Notes: This figure plots the transition path for the interest rate, wages, capital, and output after the tax reform is implemented. $t = 0$ shows the calibrated steady state. Responses of wages, output, and capital are plotted in percentage deviation from steady state; the interest rate response is plotted in differences.

households. Interestingly, the efficiency component is slightly positive, as the decrease in aggregate consumption is offset by larger leisure (Table 6, top panel). Efficiency gains also result from a better allocation of hours worked, with the distribution of hours tilted toward more productive households (Table 6, bottom panel).

4.4 Exploring the phase-out of transfers: UBI and affine plans

A positive phase-out of transfers allows to implement non-monotonic marginal $t\&T$ rates, which is a property featured in the optimal plan and is also reminiscent of the optimal

Table 6: Distribution of hours: Calibration and optimal $t\&T$ system

Mean hours	Hours worked quintile				
Calibration	0.20	0.28	0.32	0.35	0.38
Optimal	0.18	0.25	0.28	0.31	0.34
Log hours deviation	Wage quintile				
Calibration	-0.03	0.00	-0.04	0.05	0.01
Optimal	-0.18	-0.07	-0.01	0.12	0.13

Notes: The top panel shows average hours worked by quintile, sorting households by hours worked. The bottom panel reports average log hours per quintile minus the mean of log hours, sorting households per wages. Statistics are reported for the calibration and for the steady state of the optimal $t\&T$ system.

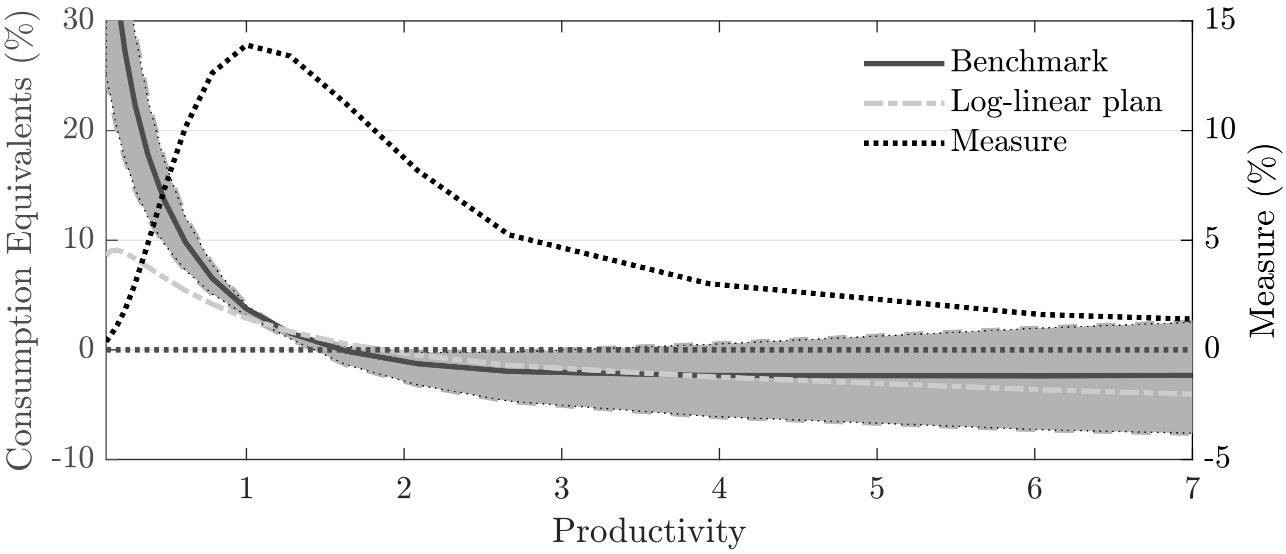

Figure 9: Optimal tax-and-transfer system: Consumption equivalent welfare gains

Notes: This figure shows welfare gains, in terms of consumption equivalent (CE), from a tax reform to the optimal system. The solid line plots the average CE by productivity level z, while the two thin dotted lines show the bottom-20% and the top-80% of the distribution of CE at each z. Both axes are cut for readability; welfare gains peak at 43% for the lowest-productivity level and converge to about -2% for top levels of productivity. The dashed line plots the average CE under the optimal log-linear plan. The right axis plots the measure of households for each productivity level.

U-shaped marginal rates often found in the public finance literature. To evaluate the importance of non-monotonic marginal rates, we compute the optimal plan when transfers do not phase out ($\xi = 0$). In this case, the shape of marginal $t\&T$ rates is monotonic and determined by the labor tax progressivity.

Eliminating the transfer phase-out allows us to discuss two common tax proposals: a UBI and an affine plan. The UBI plan eliminates the phase-out of transfers but still optimizes the progressivity of labor taxes θ. The affine plan is a UBI plan financed with flat income taxes—that is, $\theta = 0$.

The UBI plan includes large transfers at $m = 0.21$, or about \$18,700 for each household, which are optimally financed with almost flat labor taxes at $\theta = 0.04$. Transfers in the optimal UBI amount to a large government outlay representing 18% of GDP, compared to 5% in the benchmark plan. In order to finance the transfers and preserve labor supply incentives, income-tax progressivity falls so as to maintain roughly flat marginal $t\&T$ rates. Labor taxes are thus higher than in the plan with a phase-out and almost constant across households, spanning from 56% in the first income quintile to 61% in the top income quintile. Although with different tax and transfer rates, the UBI plan achieves a remarkably comparable level of redistribution as with a phase-out, with comparable average $t\&T$ rates across households (see Table 5).

The optimal affine plan is similar to the UBI, as the latter uses almost flat taxes. The affine plan features a lump-sum transfer of $20,300 to each household and a tax rate of 60%.

The welfare gain of the UBI is 5.36% in consumption-equivalent terms, while the affine gains are 5.26%, both close to the gains in the plan with a phase-out. Thus, our framework is supportive of lump-sum transfers. However, while the phasing out of transfers generates only modest additional welfare gains, it is associated with lower labor tax rates, which may be easier to implement in practice.

Overall, the UBI and affine exercises point to the importance of the intercept, as highlighted in the simple model. Transfers, even if lump sum, allow to separate the progressivity of average and marginal tax rates. The phasing-out of transfers additionally allows for non-monotonic marginal tax rates. Disentangling the progressivity of average and marginal tax rates generates most of the welfare gains. Non-monotonic marginal rates are less important for welfare.

Comparison to the optimal log-linear plan.—We compute the optimal plan using a log-linear labor tax function, as in the analytical section.[25] The optimal plan features large progressivity, at $\tau = 0.26$ when maximizing steady-state welfare and $\tau = 0.39$ when including transitions, compared to the estimated $\tau = 0.18$ for the current U.S. system. Yet, compared to the optimal plan with targeted transfers, the optimal log-linear plan achieves less redistribution and at the cost of strongly increasing marginal $t\&T$ rates (see Table 5). Welfare gains are large, at +2.88%, but smaller than with targeted transfers, in particular for households at the bottom of the income distribution (see Figure 9). Interestingly, the most productive households are also worse-off under the log-linear plan, because of the higher tax rates.

Optimal UBI with fixed income-tax progressivity.—To compare our results with those of Guner, Kaygusuz, and Ventura (2021), we compute the lump-sum transfer that optimizes steady-state welfare when fixing income-tax progressivity to its status quo level. We find an optimal transfer of $10,700, a number close to the optimal transfer of $10,400 for a couple with two children reported in Guner, Kaygusuz, and Ventura (2021). Accounting for transitions and optimizing on income-tax progressivity both increase the optimal transfer.

5 Main quantitative determinants of the optimal plan

This section quantitatively investigates the main determinants of the optimal $t\&T$ plan. We first explore the effects of changing the tails of the income distribution as well as changing the distribution of income risk. We also report alternative calibrations of public spending and preferences. To ease comparison across environments, we focus on plans

[25] See Appendix D.3 for more details.

with no phase-out ($\xi = 0$).[26]

5.1 Income distribution: A story of two tails

A key insight of the analytical model is the optimal negative relation between transfers and progressivity. This relation remains valid in the quantitative model: for a given phase out ξ, larger transfers m are optimally associated with lower income-tax progressivity θ. We refer to this relation as the transfer-progressivity locus. To shed light on the role of the income distribution, we explore how this locus shifts with the left and right tails of the income distribution.

We compute the transfer-progressivity locus in our benchmark and in another two economies: a "No Pareto" economy and a "Richer Poor" economy. For the "No Pareto" economy, we remove the Pareto tail adjustment to the highest values of the productivity grid. For the "Richer Poor" we increase the lowest 20% of the productivity distribution to equate its 20th percentile. In both cases, we readjust all remaining parameters to match the same calibration targets as in our benchmark, except for the income distribution. The left panel of Figure 10 shows the optimal transfer-progressivity locus for the benchmark and the two additional cases. The x-axis plots transfers in dollars, which are pinned down by m. The y-axis plots the difference between the marginal rates of the top 10% and of the entire distribution, a measure that closely correlates with the income-tax progressivity parameter θ. The right panel of Figure 10 reports the productivity distribution for the benchmark and the two alternative economies. Table 7 reports more statistics on the optimal plans.

In the "No Pareto" economy, the transfer-progressivity locus shifts down. That is, for a given level of transfer, the optimal progressivity falls significantly, reducing the difference in marginal rates at the top by about 13 percentage points. This finding confirms the results in Mankiw, Weinzierl, and Yagan (2009) and Heathcote and Tsujiyama (2021), which emphasize the importance of the right tail of the income distribution in determining the optimal slope of the marginal rates at the top. Remarkably, while the optimal plan features regressive income taxes—as in the simple model of Section 2—the optimal level of transfers is almost the same as in the benchmark economy. Thus, the concentration of income at the top does not significantly affect the size of optimal transfers, but it does change the optimal way of financing them. With the Pareto tail, a planner is able to raise sufficient revenues from the top and prefers lower tax rates on middle-income households.

In contrast, the transfer-progressivity locus barely shifts in the "Richer Poor" economy. That is, for a given level of transfer, the optimal level of progressivity is similar to the benchmark. However, optimal transfers decrease by about \$8,000, thus implying a larger optimal income-tax progressivity.

[26] Table D.2 in Appendix D.4 reports optimal plans with a positive phase-out.

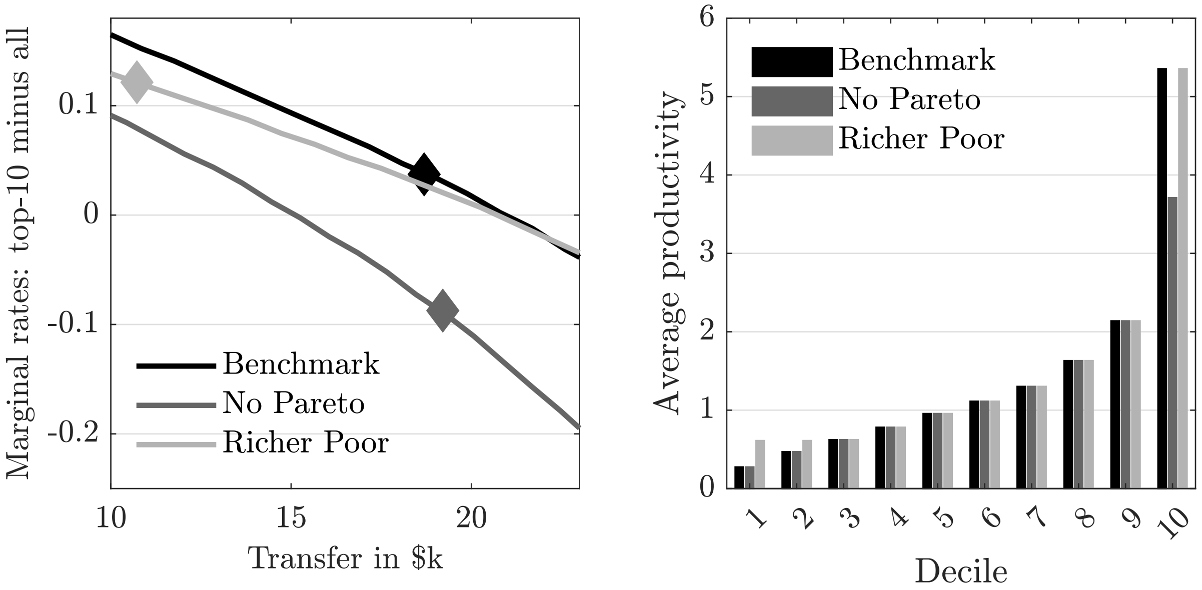

Figure 10: Optimal transfer-progressivity locus and the income distribution

Notes: The left panel plots, for each level of transfer, the optimal income-tax progressivity—proxied by the difference between the marginal rates of the top 10% and of the entire distribution. Diamonds depict the optimal transfer-progressivity pair on each transfer-progressivity locus. Results are reported for three economies: the benchmark, the "No Pareto", and the "Richer Poor" economies. The phase-out parameter ξ is fixed to zero. The right panel shows the average productivity level z by decile in each economy.

5.2 Income risk

The distribution of income risk only moderately affects the optimal fiscal plan. Table 7 reports the optimal plan when labor productivity follows an AR(1) process, abstracting from the Pareto tail for simplicity. We report two cases. First, we keep the persistence of the AR(1) fixed to the calibrated value and set its innovation variance to match the same overall productivity variance as in the GMAR case. We label this case as "No Pareto Normal" and compare it to the "No Pareto" economy, which features the GMAR process described in equation (15). Next, we increase the persistence from $\rho = 0.935$ to $\rho = 0.95$ and decrease the standard deviation of the innovation from 0.261 to 0.229 to keep the overall unconditional variance of productivity unchanged. We label this case as "No Pareto Persistent".

As compared to the GMAR, the AR process with same persistence has fewer large negative shocks but more frequent medium-sized shocks. These differences have offsetting effects on optimal taxes: while the less frequent left tail shocks reduce the need for insurance, the more frequent medium-sized shocks do the opposite. This offsetting results in a similar progressivity and somewhat larger transfers—from $19,200 in the "No Pareto" case to $19,800 in the "No Pareto Normal" case.[27] Remarkably, the optimal plan is

[27]This result is consistent with De Nardi, Fella, and Paz-Pardo (2019), which finds that the welfare

Table 7: Optimal $t\&T$ plans under alternative calibrations

	[1] Transfers	[2] Δ Marginal rate	[3] Average rate	[4] CE	[5] Fraction with CE>0
Income distribution					
Benchmark	\$18,700	4%	56%	5.36%	77%
No Pareto	\$19,200	-9%	62%	3.12%	71%
Richer Poor	\$10,700	12%	37%	2.32%	80%
Income risk					
No Pareto Normal	\$19,800	-9%	64%	3.74%	73%
No Pareto Persistent	\$20,700	-10%	66%	4.71%	72%
Fiscal Space					
Low Spending	\$21,100	4%	55%	7.58%	81%
Preferences					
Low σ & φ	\$13,900	6%	45%	2.11%	68%

Notes: This table presents statistics of the optimal $t\&T$ plan with no phase-out for the benchmark economy and 7 alternative calibrations. Column [1] reports the transfer in dollars; [2] the difference between the marginal labor income tax rates of the top 10% and of the entire distribution; [3] the average labor income tax rate, averaged over the entire distribution; [4] the welfare gain in consumption equivalent terms; and [5] the fraction of the population supporting the reform.

largely unchanged with more persistent income shocks, once the variance of the process is recalibrated. Optimal transfers increase mildly in the "No Pareto Persistent" case, at \$20,700. Overall, the distribution of income risk does not drastically change the optimal $t\&T$ plan.

5.3 Fiscal space

The optimal plan is sensitive to the level of spending G. Table 7 reports the optimal plan in an alternative "Low Spending" calibration. In the initial steady state, we decrease λ, the level parameter of labor income taxes, to target a spending-over-output ratio of 16%—that is, 4 percentage points lower than in the benchmark calibration. We then recalibrate the remaining parameters to match the same calibration targets as in the benchmark.

As expected, with more fiscal space, optimal transfers increase, from \$18,700 to \$21,100. Welfare gains are also larger. These results are consistent with the findings in Heathcote and Tsujiyama (2021).

cost of idiosyncratic risk is somewhat lower when modeling higher-order moments of income risk.

5.4 Preferences

The optimal plan is sensitive to preference parameters. Table 7 reports the optimal plan in an alternative calibration with lower risk aversion and higher Frisch elasticity, which we label as "Low σ & φ". We set the risk aversion to $\sigma = 1.5$ and find the Frisch elasticity $\varphi^{-1} = 0.5$ such that wealth effects remain unchanged. We then readjust all remaining parameters, including the income distribution, to match the same calibration targets as in the benchmark.

With lower risk aversion and higher Frisch elasticity, optimal transfers decrease to \$13,900. This result is intuitive: with higher Frisch elasticities, the distortionary cost of labor taxes increases, while, with lower risk aversion, the insurance/redistribution gains of transfers decrease. The welfare gains of the optimal plan are also scaled down, to 2.11%. Still, a large majority of households would support the reform (68%).

To summarize, the comparative statics of this section provide three takeaways. First, the left tail of the income distribution determines optimal transfers, whereas the right tail determines optimal income-tax progressivity. Second, the distribution of income risk does not drastically change the optimal $t\&T$ plan. Third, the optimal plan provides less redistribution/insurance with higher Frisch elasticity or lower risk aversion. Yet in all cases, optimal transfers are larger than currently in the U.S., welfare gains of a tax reform are sizable, and a vast majority of households would support it.

6 Conclusion

In this paper, we studied the optimal design of the tax-and-transfer system. We developed a tractable analytical model demonstrating an optimally negative relation between transfers and income-tax progressivity, and showed that adding a transfer to a log-linear tax induces welfare gains almost as large as in the second-best allocation. We then quantified the optimal fiscal plan in a rich dynamic model calibrated to the U.S. economy. We found that optimal transfers should be generous, with a slow phase-out, and financed with moderate income-tax progressivity. That is, the optimal plan features more progressive average than marginal $t\&T$ rates. Furthermore, there are large welfare gains from implementing this plan, which would be supported by a majority of households.

Our Ramsey approach is suitable for a quantitative and dynamic evaluation of efficiency and redistribution concerns. The instruments we used are simple and intuitive, and they resemble current policies implemented by many countries. Yet they are richer than what is typically used in the Ramsey literature, and flexible enough to generate nonlinear non-monotonic $t\&T$ schedules. As such, our analysis contributes to bridging the gap between Mirrlees and Ramsey traditions.

A Analytical model

A.1 Optimal plan without transfers

This section presents the derivation of the welfare formula when the government uses a log-linear income tax and no lump-sum transfer.

Household problem.—Define $\alpha_i \equiv \log z_i$, with $\alpha_i \sim \mathcal{N}(-v_\omega/2, v_\omega)$ so that $\mathbb{E}[z_i] = 1$. The consumers solve a static problem: they maximize utility (1) given the budget constraint (2), with $T = 0$. The first-order conditions deliver (5).

Computing λ.—Output is given by $Y = \int y_i d_i = \int \exp(\alpha_i) n_i d_i = n_0$. To compute λ, we also need to compute $\tilde{Y} \equiv \int y_i^{1-\tau} d_i = \int [n_0 \exp(\alpha_i)]^{1-\tau} d_i = n_0^{1-\tau} \exp\left(-\tau(1-\tau)\frac{v_\omega}{2}\right)$. Therefore, using the government's budget constraint (3), we can express λ as

$$\lambda = \frac{Y - G}{\tilde{Y}} = \frac{n_0 - G}{n_0^{1-\tau}} \exp\left(\tau(1-\tau)\frac{v_\omega}{2}\right) \equiv \lambda_0. \tag{17}$$

Welfare.—To compute welfare in closed form, we plug the equilibrium values of consumption and hours worked into the utility function and integrate over the distribution of households, we obtain

$$W(\tau) = \int u_i d_i = \log \lambda + \frac{1-\tau}{1+\varphi} \log\left(\frac{1-\tau}{B}\right) - (1-\tau)\frac{v_\omega}{2} - \frac{1-\tau}{1+\varphi},$$

and, using the closed-form solution (17) for λ, we obtain equation (6) in the main text.

A.2 Optimal plan with transfers

This section presents the case where the government is endowed with a log-linear income tax and a lump-sum transfer.

A.2.1 Representative agent

We first abstract from heterogeneity and assume $v_\omega = 0$. We show that: (1) for any level of transfer there exists a progressivity that implements the efficient allocation; and (2) this progressivity is decreasing in the transfer.

First-best.—Let $n^\star(G)$ denote the first-best allocation, which maximizes the utility of the representative agent given the resource constraint (4). It is characterized by

$$B\,(n^\star(G))^\varphi\,(n^\star(G) - G) = 1. \tag{18}$$

Second-best.—We turn to the optimal allocation when spending is financed with a log-linear income tax and a transfer. The household chooses $\{c, n\}$ to maximize utility (1) given the budget constraint (2). The household first-order condition is

$$Bn^\varphi = \frac{\lambda\,(1-\tau)\,n^{-\tau}}{\lambda n^{1-\tau} + T}.$$

Efficient second-best.—Rearranging and using the government's budget constraint (3), it follows that for any level of transfer, the efficient allocation $n^\star(G)$, characterized by (18), can be implemented with progressivity $\tau^\star(G,T)$ equal to

$$\tau^\star(G,T) = -\frac{G+T}{n^\star(G)-(G+T)}. \tag{19}$$

Furthermore, (19) implies that the optimal progressivity τ decreases with transfers T.

To understand why the first-best allocation can be implemented for *any* level of transfers, we first focus on the case when $T=0$. The labor policy function, characterized in (5), shows a one-to-one negative relationship between n and τ, which ensures that a planner can always pick the progressivity τ to implement $n^\star(G)$. When transfers are not zero, the policy function for labor does not admit a closed-form solution. However, labor supply decreases with transfers T because of a wealth effect. This property explains the optimal negative relationship between τ and T obtained in (19). Progressivity falls as transfers increase, such that labor supply remains at $n^\star(G)$ despite the more generous transfers.[28] At the optimal labor level, for any T, the marginal rate is zero.

Comparison with the "fiscal pressure" effect.—Note that the effect of transfers on progressivity is qualitatively different from the one of spending. Larger spending increases the first-best labor supply, but only taxes directly impact the household labor policy. With higher spending, optimal progressivity decreases so that labor supply reaches its new first-best level. Transfers, by contrast, do not alter the first-best labor supply but weaken the household's incentives to provide labor. With higher transfers, optimal progressivity decreases so that labor supply remains at its unchanged first-best level.

Figure A.1 presents a numerical illustration of the relationship between transfers and progressivity for the representative agent case. As can be seen, optimal progressivity τ declines as transfers T increase. We also plot the optimal relation between τ and T for higher spending: a larger G generates a new curve, associated with a new level of labor, while a larger T moves the optimum on the existing curve.

A.2.2 Heterogeneity

We now derive welfare as a function of progressivity τ and the transfer T when $v_\omega > 0$. The logic of the derivation is the same as with no transfers. However, we cannot express the policy function for labor in closed form, so we linearize around $T=0$.

Household problem.—The first-order condition of the household problem reads

$$Bn_i^{1+\varphi} + \frac{T}{\lambda} Bn_i^{\varphi+\tau} \exp(-(1-\tau)\alpha_i) - (1-\tau) = 0. \tag{20}$$

Equation (20) defines the function $F(T,\lambda,n_i)$ s.t. $F(T,\lambda,n_i)=0$. At the optimum, the labor

[28] Equation (19) retrieves an old result in the Ramsey literature. When $\tau=0$, the optimal transfer is $T=-G$. That is, the government finances all spending with a lump-sum tax and sets the distortionary flat tax rate at zero.

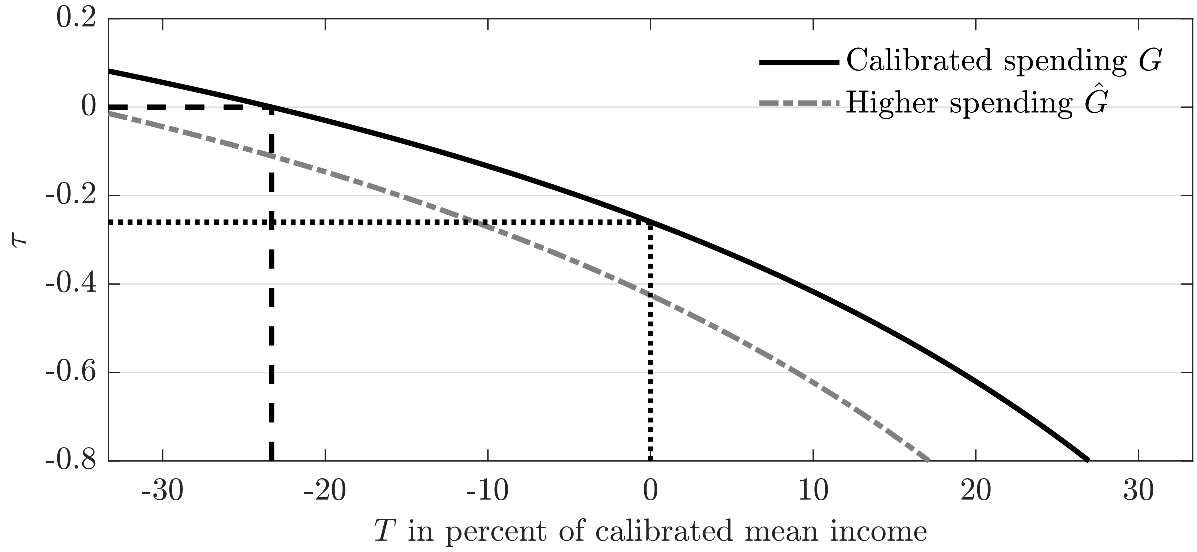

Figure A.1: Representative agent: implementing the first-best allocation

Notes: This figure shows the combinations of lump-sum transfer T and income-tax progressivity τ that implement the first-best allocation in the representative-agent case. The dashed line highlights the case of zero transfer, in which progressivity is negative: $\tau = -0.26$. The dotted line marks the scenario of an affine tax system, $\tau = 0$, in which a lump-sum tax finances all government spending: $T = -G$. The dash-dotted line shows the optimal combinations of T and τ for larger spending, at $\hat{G} = 1.5G$.

decision is such that, for a given T, $F(T, \lambda, n_i(T, \lambda)) = 0$.

Linear approximation of labor policy.—At $T = 0 \equiv T_0$, $n_i(T_0, .) = n_0$. The implicit function theorem applies, and we can compute the slope of n_i in the neighborhood of T_0 as

$$\left.\frac{\partial n_i(T, \lambda)}{\partial T}\right|_{(T_0, n_0, \lambda_0)} = -\frac{\left.\frac{\partial F(T, \lambda, n_i)}{\partial T}\right|_{(T_0, n_0, \lambda_0)}}{\left.\frac{\partial F(T, \lambda, n_i)}{\partial n_i}\right|_{(T_0, n_0, \lambda_0)}}.$$

Let $\eta \equiv \exp\left[(1-\tau)\, v_\omega\right]$, where we omit the dependence of η on τ and v_ω to ease notation. We compute the two partial derivatives and obtain a linear approximation around (T_0, n_0, λ_0) of $n_i(T)$ denoted $\hat{n}_i(T)$:

$$\hat{n}_i(T) = n_0 + T \left.\frac{\partial n_i(T)}{\partial T}\right|_{(T_0, n_0, \lambda_0)} = n_0 - \frac{T}{1+\varphi} \frac{n_0}{n_0 - G} \frac{\eta^{-\frac{\tau}{2}}}{\exp[(1-\tau)\alpha_i]}, \tag{21}$$

which delivers equation (7).

Computing λ.—Again, we need to compute Y and $\tilde{Y}$, which we linearize as $\hat{Y}$ and $\hat{\tilde{Y}}$. We start with output:

$$\hat{Y} = \int y_i d_i = \int \exp\left(\alpha_i\right) n_i d_i = n_0 - \frac{T}{1+\varphi} \frac{n_0}{n_0 - G} \eta^{-\tau}. \tag{22}$$

To obtain $\tilde{Y}$, we first approximate $\hat{n}_i^{1-\tau}$. Using equation (21) and linearizing, we get

$$\hat{n}_i^{1-\tau}(T) = \left[n_0 - \frac{T}{1+\varphi}\frac{n_0}{n_0-G}\frac{\eta^{-\frac{\tau}{2}}}{\exp[(1-\tau)\alpha_i]}\right]^{1-\tau} = n_0^{1-\tau} - \frac{T}{1+\varphi}\frac{n_0^{1-\tau}}{n_0-G}\frac{\eta^{-\frac{\tau}{2}}(1-\tau)}{\exp[(1-\tau)\alpha_i]}.$$

It follows that

$$\hat{\tilde{Y}} = \int\left[n_0^{1-\tau} - \frac{T}{1+\varphi}\frac{n_0^{1-\tau}}{n_0-G}\frac{\eta^{-\frac{\tau}{2}}(1-\tau)}{\exp[(1-\tau)\alpha_i]}\right]\exp\left[(1-\tau)\alpha_i\right]d_i = n_0^{1-\tau}\eta^{-\frac{\tau}{2}}\left[1 - \frac{T}{1+\varphi}\frac{1-\tau}{n_o-G}\right].$$

Using these expressions, we can compute λ using the government budget constraint (3), and obtain, taking derivatives and linearizing around $T = 0$,

$$\hat{\lambda}(T) = \lambda_0 + \frac{T}{1+\varphi}\frac{1}{\eta^{-\frac{\tau}{2}}n_0^{1-\tau}}\left[-\frac{n_0}{n_0-G}\eta^{-\tau} - (\varphi+\tau)\right]. \tag{23}$$

Welfare.—We approximate utility around $T = 0$. The utility of an agent is given by

$$u_i = \log\left[\lambda\left[\exp\left(\alpha_i\right)n_i\right]^{1-\tau} + T\right] - \frac{B}{1+\varphi}n_i^{1+\varphi},$$

which, using our expressions (23) for $\hat{\lambda}(T)$ and (21) for $\hat{n}_i(T)$, can be approximated as

$$\begin{aligned}\hat{u}_i &= u_{i,0} + T\left\{\frac{\lambda'(T)\exp\left[(1-\tau)\alpha_i\right]n_0^{1-\tau} - \lambda_0\frac{1-\tau}{1+\varphi}\frac{n_0^{1-\tau}}{n_0-G}\eta^{-\frac{\tau}{2}} + 1}{\lambda_0\exp\left[(1-\tau)\alpha_i\right]n_0^{1-\tau}} + \frac{B}{1+\varphi}\frac{n_0^{1+\varphi}}{n_0-G}\frac{\eta^{-\frac{\tau}{2}}}{\exp[(1-\tau)\alpha_i]}\right\} \\ &= u_{i,0} + T\left\{\frac{1}{1+\varphi}\frac{1}{n_0-G}\left[-\frac{n_0}{n_0-G}\eta^{-\tau} - (\varphi+\tau)\right] - \frac{1-\tau}{1+\varphi}\frac{1}{n_0-G}\frac{\eta^{-\frac{\tau}{2}}}{\exp[(1-\tau)\alpha_i]}\cdots\right. \\ &\qquad \left. + \frac{1}{n_0-G}\frac{\eta^{-\frac{\tau}{2}}}{\exp[(1-\tau)\alpha_i]} + \frac{B}{1+\varphi}\frac{n_0^{1+\varphi}}{n_0-G}\frac{\eta^{-\frac{\tau}{2}}}{\exp[(1-\tau)\alpha_i]}\right\}.\end{aligned}$$

Integrating this equation yields $W(\tau, T) = W(\tau, 0) + \hat{\Omega}(\tau, v_\omega)T$, with $\hat{\Omega}(\tau, v_\omega)$ defined as

$$\hat{\Omega}(\tau, v_\omega) \equiv \frac{1}{1+\varphi}\frac{1}{n_0-G}\left[-\frac{n_0}{n_0-G} + 1 - \tau\right] + \frac{1}{1+\varphi}\frac{n_0}{(n_0-G)^2}\left[-\eta^{-\tau}+1\right] + \frac{1}{n_0-G}\left[\eta^{1-\tau}-1\right], \tag{24}$$

where, again, we omit the dependence of n_0 on τ and of η on τ and v_ω.

Welfare decomposition: efficiency.—The first term in equation (24), which equates $\Omega_e(\tau, 0)$ as defined in (8.a), can be rearranged as

$$\Omega_e(\tau, 0) = -\frac{1}{n_0-G}\frac{1}{1+\varphi}\frac{n_0}{n_0-G} + \frac{1-\tau}{n_0}\frac{1}{1+\varphi}\frac{n_0}{n_0-G}.$$

It is equal to the marginal utility of aggregate consumption with no transfer $u_c(C_0) = 1/(n_0 - G)$, multiplied by $\partial\hat{Y}^{ra}(T)/\partial T$, where $\hat{Y}^{ra}(T)$ is defined as the representative-agent version of (22) (i.e. with $\eta = 1$), evaluated at $T = 0$; plus the marginal utility of aggregate leisure $-u_n(n_0) = Bn_0^{\varphi} = (1-\tau)/n_0$, multiplied by $\partial\hat{n}^{ra}(T)/\partial T$, where $\hat{n}^{ra}(T)$ is defined as the representative-agent version of (21), evaluated at $T = 0$.

Adding the second term in equation (24) and rearranging, we retrieve $\Omega_e(\tau, v_\omega)$ in (8.b)

$$\Omega_e(\tau, v_\omega) = \Omega_e(\tau, 0) + \frac{1}{n_0 - G}\left[-\frac{1}{1+\varphi}\frac{n_0}{n_0 - G}\eta^{-\tau} + \frac{1}{1+\varphi}\frac{n_0}{n_0 - G}\right],$$

where the additional parenthesis, which captures the heterogeneity of wealth effects on labor supply, is equal to the marginal utility of aggregate consumption $u_c(C_0)$, multiplied by $\partial\hat{Y}(T)/\partial T$ minus $\partial\hat{Y}^{ra}(T)/\partial T$. Approximating further using $\exp(x) \approx 1 + x$ delivers

$$\Omega_e(\tau, v_\omega) = \Omega_e(\tau, 0) + \frac{\tau(1-\tau)}{n_0 - G}\frac{n_0}{n_0 - G}\frac{v_\omega}{1+\varphi},$$

which shows that the additional term changes non-monotonically with τ. Transfers reduce labor supply more in the presence of heterogeneity, as the larger wealth effect of the poor more than offsets the smaller wealth effect of the rich. Higher progressivity reduces this dispersion of wealth effects, bringing labor supply closer to the representative agent case. However, the effect on output depends not only on labor supply but also on the distribution of z, which makes the effect of τ non-monotonic. In our calibrations, this second term is typically small, and $\partial[\Omega_e(\tau, v_\omega)]/\partial\tau < 0$ holds quantitatively, so that the total efficiency gains of transfers decrease with progressivity.

Welfare decomposition: redistribution.—The third term in equation (24) can be rewritten as

$$\Omega_r(\tau, v_\omega) = \frac{1}{n_0 - G}\eta^{1-\tau} - \frac{1}{n_0 - G}$$

and is equal to the average marginal utility across households, $\int u_c(\lambda_0 (z_i n_0)^{1-\tau}) d_i$, minus marginal utility of aggregate consumption $u_c(C_0)$. Approximating further using $\exp(x) \approx 1 + x$ delivers equation (8.c).

Derivatives.—Claim 1 states that $\Omega_e(\tau, 0)$ is decreasing in τ. To show that, we sign the derivative equal to

$$\frac{1}{1+\varphi}\frac{1}{n_0 - G}\left(\frac{n_0}{n_0 - G}\frac{1}{1+\varphi}\left[-\frac{n_0 + G}{n_0 - G}\frac{1}{1-\tau} + 1\right] - 1\right).$$

When $\tau > 0$, $\frac{1}{1-\tau} > 1$ and $\frac{n+G}{n-G} \geq 1$ such that $-\frac{1}{1-\tau}\frac{n+G}{n-G} + 1 < 0$ and the derivative is always negative. More generally, we need to show that the parenthesis is negative. As $\varphi > 0$, it is sufficient to show that

$$\frac{n_0}{n_0 - G}\left[-\frac{n_0 + G}{n_0 - G}\frac{1}{1-\tau} + 1\right] \leq 1$$

is always true. This condition can be rewritten as

$$-(1-\tau)G^2 \leq n_0(n_0 + \tau G). \tag{25}$$

Equation (25) always holds as the left-hand side is negative $\forall\tau$, while the right-hand side is positive as $\tau \geq -1$ and $G \leq n_0$ by feasibility.

Claim 2 states that $\Omega_r(\tau, v_\omega)$ is decreasing in τ. The derivative reads

$$(1-\tau)v_\omega \frac{1}{n_0(\tau) - G}\left(-2 + \frac{1}{1+\varphi}\frac{n_0(\tau)}{n_0(\tau) - G}\right), \tag{26}$$

which is negative $\forall \tau \in [-1; 1]$ when $G = 0$. When $G > 0$, equation (26) is negative on $\tau \in [-1, \hat{\tau}(G)]$, with $\hat{\tau}(G) \equiv 1 - BG^{1+\varphi}\left(\frac{2(1+\varphi)}{2(1+\varphi)-1}\right)^{1+\varphi}$, which is above 0.99 in our calibration.

A.3 Mirrlees allocation

This section briefly presents the Mirrlees problem.[29]

Household.—The utility of a household with productivity z and facing a tax schedule $\Xi(.)$ can be written as

$$v(z, \Xi(.)) \equiv \max_{c,n} \log c - B\frac{n^{1+\varphi}}{1+\varphi} \text{ s.t. } c = zn - \Xi(zn).$$

Let $\nu(z, \Xi(.))$ denote the labor policy.

Government.—The government chooses the taxation schedule $\Xi(.)$ to maximize the sum of utilities across households subject to its budget constraint:

$$\max_{\Xi(.)} \int_z v(z, \Xi(.))dF_z(z) \text{ s.t. } \int_z \Xi(z\nu(z, \Xi(.)))dF_z(z) \geq G.$$

where F_z is the distribution of households' productivity.

Note that the level of labor disutility B is irrelevant to the problem, as long as G is recalibrated so that G/Y remains constant. Indeed, the parameter B simply scales the entire economy up and down and does not affect the efficiency-redistribution trade-off under homothetic preferences.[30]

Heathcote and Tsujiyama (2021) presents a richer environment with individual productivity

[29] Alternatively, the Mirrlees problem can be written as a mechanism design problem. The utility of a household with true productivity z and reported productivity $\tilde{z}$ is given by

$$U(z, \tilde{z}) \equiv \log c(\tilde{z}) - \frac{B}{1+\varphi}\left(\frac{y(\tilde{z})}{z}\right)^{1+\varphi}$$

where $c(\tilde{z})$ and $y(\tilde{z})$ denote consumption and labor income. Given the productivity distribution denoted F_z, the government chooses allocations $\{c(z), y(z)\}_z$ to maximize $\int_z U(z, z)dF_z(z)$ subject to the resource constraint and a set of incentive constraints:

$$\int_z c(z)dF_z(z) + G = \int y(z)dF_z(z), \text{ and } U(z, z) \geq U(z, \tilde{z}) \;\forall (z, \tilde{z}).$$

See Heathcote and Tsujiyama (2021) for more details.

[30] Consider an alternative labor disutility parameter $\hat{B}$, and a productivity process rescaled by $\xi \equiv (\hat{B}/B)^{1/(1+\varphi)}$. Then

$$\log c(\tilde{z}) - \frac{\hat{B}}{1+\varphi}\left(\frac{y(\tilde{z})}{\xi z}\right)^{1+\varphi} = \log c(\tilde{z}) - \frac{B}{1+\varphi}\left(\frac{y(\tilde{z})}{z}\right)^{1+\varphi}.$$

shocks which are insurable within families, in addition to the standard non-insurable shocks common to all individuals in a family. As income is taxed at the family level, insurable shocks only map to a different labor disutility parameter B of the family/household (see equation (9) of their paper). Thus, our setup is equivalent to theirs, and we retrieve their welfare numbers when using their calibration for non-insurable shocks (see Table 1), despite the fact that our model does not feature insurable shocks.

B Data

B.1 Current Population Survey

We use the CPS to measure taxes and transfers as well as hours worked. The main data source is the CPS Annual Social and Economic Supplement (ASEC) for 2013. We start from the CPS version provided by IPUMS (Flood et al. 2021). Tax variables are imputed using the Census Bureau's CPS ASEC Tax Model (Lin 2022). Transfers are reported by households, but they are severely underreported. We follow CBO procedure (Habib 2018), and impute transfers when possible.[31] Table B.1 contains a list of all tax-and-transfer variables that we use in the benchmark analysis and in various robustness checks.

Table B.2 compares the aggregate amounts of transfer and tax credit variables in the CPS to their counterparts from national accounts. The aggregate data is from the IRS Statistics of Income for tax credits, from NIPA (Table 3.12) for SNAP, SSI, and Medicaid, and the Congressional Research Service for housing assistance (public housing and rental assistance).[32] We do not report a data counterpart for welfare, as this measure contains a number of small programs. The fit is generally good; tax credits are underestimated, which is a common problem when using survey data only to impute credits because of the difficulty of identifying qualifying dependents (Lin 2022; Meyer et al. 2020).

B.1.1 Benchmark tax-and-transfer estimates

We first define measures of income, taxes, and transfers. All variables are aggregated at the household level.

Our benchmark measure of labor income includes wage and salary income (incwage), non-farm business income (incbus), farm income (incfarm), unemployment benefits (incunemp), and

[31] The Stata codes implementing the CBO imputation procedure are provided at https://github.com/US-CBO/means_tested_transfer_imputations, which we last downloaded on December 8, 2022. The merge between CPS and CBO imputed data requires original CPS identifiers, which are not included in IPUMS. We merge the original identifiers to the IPUMS version of the CPS and use those to combine it with the CBO imputation. This sequential merge procedure is described here: https://blog.popdata.org/mergecpsfile/. We obtain the original CPS data from the Census at https://www.census.gov/data/datasets/time-series/demo/cps/cps-asec.2013.html, last downloaded on December 8, 2022.

[32] For the housing assistance measure we use the CBO imputation procedure. As housing assistance, the CBO includes public housing and rental assistance; see https://www.cbo.gov/system/files/2021-09/57460-Transfers.pdf. We obtain an estimate for the sum of these two components from the Congressional Research Service; see https://crsreports.congress.gov/product/pdf/RL/RL34591.

Table B.1: Tax-and-transfer variables

Variable name	**Interpretation**	**Source**
fedtax	Federal income tax before credits	IPUMS CPS
fedtaxac	Federal income tax after credits	IPUMS CPS
ctccrd	Child tax credit (CTC)	IPUMS CPS
actccrd	Additional child tax credit (ACTC)	IPUMS CPS
eitcred	Earned income tax credit (EITC)	IPUMS CPS
fica	Payroll taxes	IPUMS CPS
fedretir	Federal retirement payroll deduction	IPUMS CPS
statetax	State income tax before credits	IPUMS CPS
stataxac	State income tax after credits	IPUMS CPS
snap_impute_val	SNAP (food stamps)	CBO imputed
ssi_impute_val	Supplemental security income (SSI)	CBO imputed
housing_assist_impute_val	Housing assistance	CBO imputed
mcaid_impute_val	Medicaid	CBO imputed
incwelfr	Various public assistance programs	IPUMS CPS

Notes: This table summarizes the tax-and-transfer variables used for benchmark estimates and robustness analysis. As federal income tax before credits includes non-refundable credits, the CTC is accounted for in fedtax. The difference between federal income tax before and after credits corresponds to refundable tax credits, that is, the EITC and the ACTC.

Table B.2: Transfer programs: comparison with national accounts

Program	**Aggregate data**	**CPS**
SNAP	74.9	75.2
Housing Assistance	37.3	36.0
Welfare	-	5.6
SSI	53.1	52.1
Medicaid	417.5	310.0
EITC	65.4	45.0
CTC	27.9	30.5
ACTC	28.1	15.7

Notes: This table compares transfer program aggregates from the CPS to national accounts, in billions of dollars.

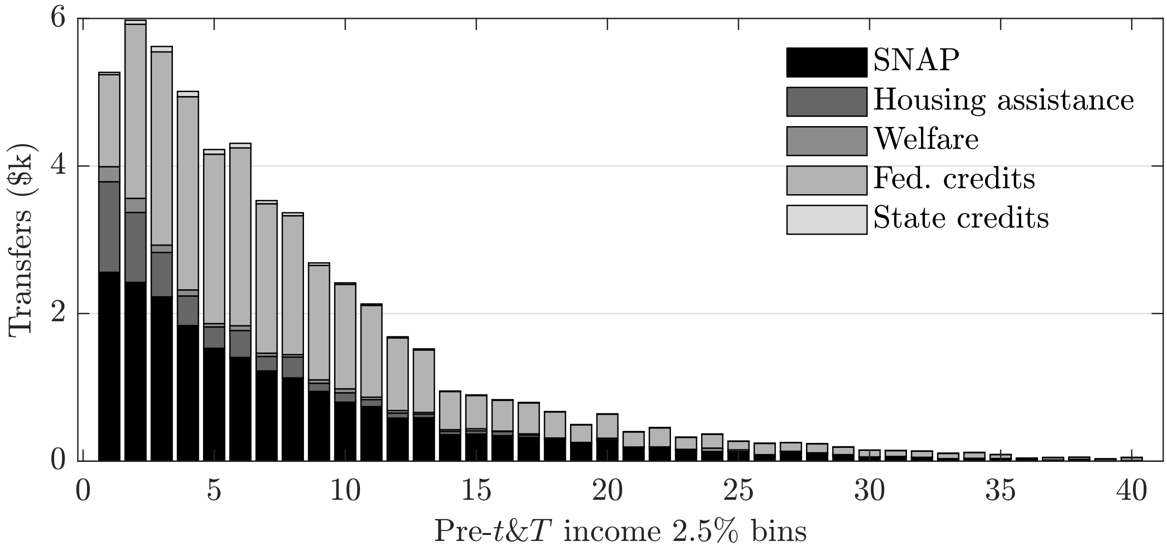

Figure B.1: Components of transfers

Notes: Components of transfers by 2.5% bins of pre-$t\&T$ income distribution.

worker's compensation (incwkom). Total income consists of labor income plus interest income (incint), dividend income (incdivid), rental income (incrent), educational assistance (inceduc), other income (incother), and averages to $\bar{y} = \$90{,}185$. Labor taxes include federal and state income taxes net of non-refunded income tax credits, as well as employee and (imputed) employer payroll taxes and federal retirement payroll deductions.[33] As we account for dividend taxes in our measure of capital taxes, we reduce taxes paid by 13.3% times dividend income.[34] Transfers include SNAP, housing assistance, welfare, and refunded federal and state income tax credits. Sample selection is as follows: (1) the head of household is between 25 and 60; (2) household total and labor pre-tax incomes are above \$5,000; (3) after tax-and-transfer total income is positive; (4) before-credit labor taxes are lower than labor income.

Figure B.1 reports the different programs along the income distribution. The larger transfers are SNAP and federal tax credits. The latter also phase out relatively slowly: they are larger at the bottom but remain significant also higher up in the income distribution.[35]

Estimates.—Table B.3 reports parameter estimates of tax-and-transfer functions. The first row presents our benchmark estimate, where we use a nonlinear estimate for the tax and the transfer functions separately. The second row estimates the parameters jointly, using only pre-tax and after-tax-and-transfer total income. The third row estimates the log-linear function,

[33] We also add imputed employer payroll taxes to labor income. We make a conservative estimate for the employer part of payroll taxes by applying rates and contribution limits to wage and salary income, but adjusting downwards if the Census imputed employee payroll tax is lower, adjusted for differences in rates.

[34] See Section B.4 for more details on estimates of the dividend tax.

[35] It may be surprising that there is any receipt of SNAP in the middle of the income distribution, as those income levels exceed the eligibility thresholds. As explained in Habib (2018), the CBO imputation procedure uses self-reported recipiency status from the CPS, where some high-income households do report receiving SNAP transfers. It may happen because SNAP eligibility is based on monthly income; some individuals can also qualify for SNAP even if they live in a high-income household, as the definition of a household in the CPS does not coincide with what is considered to determine benefit eligibility.

Table B.3: Optimal $t\&T$ plans with phasing-out under alternative calibrations

	Parameters				*Deviations*	
New functions	λ	θ	m	ξ	In levels	In logs
Benchmark	0.25	0.08	0.09	4.22	7543	0.158
Joint estimates	0.26	0.07	0.07	1.56	7510	0.156
Log-linear taxes	$1-\lambda$	τ	T			
In log	0.22	0.18	-		12677	0.165
In levels	0.22	0.09	-		7641	0.184
With a transfer	0.27	0.06	4492		7450	0.156

Notes: This table summarizes estimates of different tax-and-transfer functions.

using OLS in logs. The fourth row estimates the log-linear function in levels, whereas the last row estimates the log-linear function plus a lump-sum component, in levels as well. To measure the fit to the after-tax-and-transfer income, the last two columns report the square root of the mean squared deviation, in levels and in logs. Figure B.2 reports differences in actual minus predicted $t\&T$ paid, in levels and relative to pre-$t\&T$ income.

B.1.2 Tax-and-transfer measures: robustness

Table B.4 presents tax-and-transfer rates under alternative variables for income, taxes, and transfers, as well as the corresponding estimates of the tax functions.

First, we check robustness of our estimates with respect to the treatment of tax credits. Our benchmark measure considers refunded tax credits as transfers, while non-refunded tax credits are included in taxes. We consider two alternatives. First, we treat as transfers all refundable tax credits reported in the CPS—that is, the EITC and the ACTC—even if they are not actually refunded to households with high enough tax liabilities.[36] Second, we treat as transfers all credits reported in the CPS—that is, refundable credits plus the CTC. Most refundable credits are actually refunded: the ACTC is fully refunded by definition and the vast majority of the EITC is refunded as well. Thus, considering refunded or refundable tax credits as transfers makes no significant difference. Treating all tax credits as transfers slightly decreases the estimate of the phasing-out parameter ξ, because the CTC is received by many middle-income households, but overall our estimates are robust to the treatment of tax credits.

Next, we check robustness with respect to the treatment of unemployment insurance (UI). Our benchmark measure considers UI as labor income, because we do not model unemployment explicitly, and receiving unemployment benefits is contingent on prior work history and income. Treating UI as transfers slightly increases transfer rates at the bottom of the income distribution, and the estimated m moves up, from 0.09 in the benchmark to 0.10.

We also consider alternative definitions of income. First, our benchmark excludes income

[36] For state credits, the data does not allow to make that distinction. Also in the "refundable" case, we use "refunded" state credits.

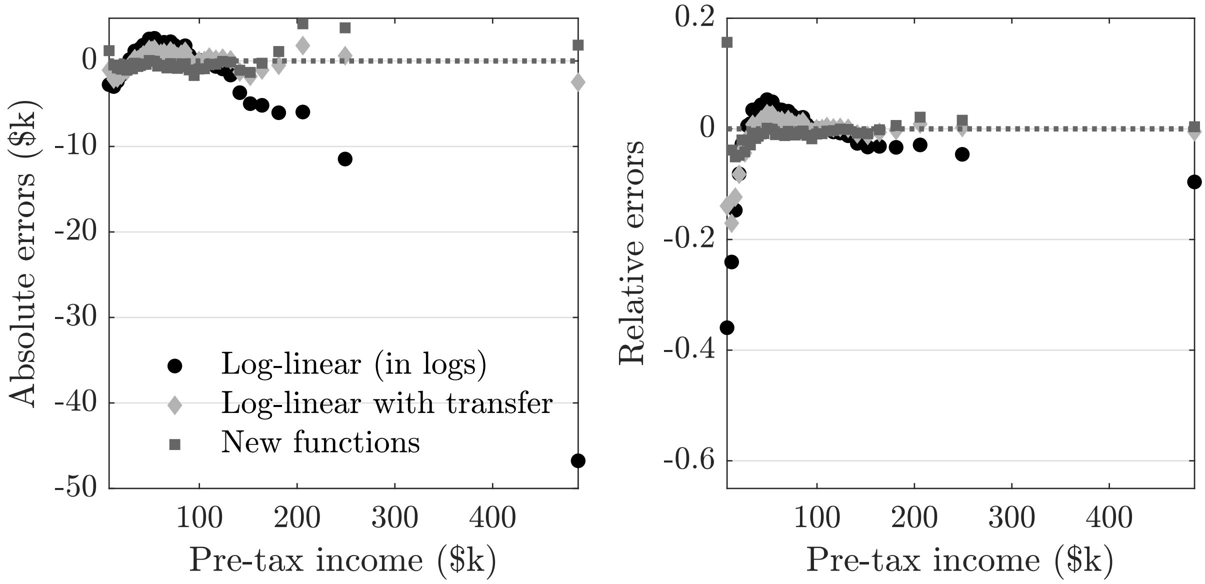

Figure B.2: Estimated $t\&T$ system: prediction errors

Notes: Prediction errors in taxes paid minus transfers received, data minus prediction. Errors are reported for three sets of estimates: (1) the log-linear income-tax function without a transfer, estimated on log income; (2) the log-linear income-tax function with a transfer, estimated on income in levels; and (3) the new tax and transfer functions. The left panel reports absolute errors in dollars, while the right panel reports relative errors as a fraction of pre-$t\&T$ income, by 2.5% bins of pre-$t\&T$ income distribution.

related to retirement, disability, and family support, all of which are related to choices and risks which are not modeled. Our estimates are robust to a broader income definition, which also includes: social security income (incss), retirement income (incretir), income from veteran's benefits (incvet), income from survivor's benefits (incsurv), income from disability benefits (incdisab), income from child support (incchild), income from alimony (incalim), and income from assistance from friends and relatives not living in the same household (incasist). Second, our benchmark allocates all of business income to labor income, because taxes on business income cannot be isolated from other taxes in the CPS. Allocating only two thirds of business income to labor income, as sometimes done in the literature, does not change our estimates.

Regarding taxes, we propose two robustness checks: on payroll and dividend taxes. First, in the benchmark we impute the share of payroll taxes paid by employers, and add it to, both, labor income and labor taxes. Ignoring the employer payroll tax decreases tax rates, by about 6% in the first quintile and about 3% in the top quintile. Accordingly, the estimate for the level of taxes, λ, decreases from 0.25 to 0.20. Second, we deduct estimated dividend taxes from the benchmark measure of taxes; not doing so leaves estimates unchanged.

The last robustness check applies to transfers. The benchmark measure for transfers excludes Medicaid and SSI, as our model abstracts from the risks insured through these programs, health and disability. Including these transfers increases significantly transfer rates, especially at the bottom of the income distribution. Accordingly, the estimated m increases significantly, from 0.09 to 0.18.

Table B.4: Tax-and-transfer rates: robustness

	Income quintiles					*Fiscal estimates*			
	Q1	Q2	Q3	Q4	Q5	λ	θ	m	ξ
Benchmark									
Transfer rates	23%	4%	1%	0%	0%				
Tax rates	16%	19%	22%	24%	29%				
Estimates						0.25	0.08	0.09	4.22
Refundable credits in transfers									
Transfer rates	23%	4%	1%	0%	0%				
Tax rates	16%	20%	22%	24%	29%				
Estimates						0.25	0.08	0.09	4.17
All credits in transfers									
Transfer rates	24%	5%	2%	1%	0%				
Tax rates	16%	21%	23%	25%	30%				
Estimates						0.25	0.07	0.07	2.88
UI in transfers rather than income									
Transfer rates	26%	5%	1%	1%	0%				
Tax rates	16%	20%	22%	24%	29%				
Estimates						0.25	0.08	0.10	3.75
Broad income definition									
Transfer rates	21%	3%	1%	0%	0%				
Tax rates	15%	18%	21%	23%	29%				
Estimates						0.25	0.08	0.09	4.35
Alternative business income allocation									
Transfer rates	23%	4%	1%	0%	0%				
Tax rates	16%	19%	22%	24%	29%				
Estimates						0.25	0.08	0.09	4.28
No employer payroll tax imputation									
Transfer rates	24%	4%	1%	0%	0%				
Tax rates	10%	14%	16%	19%	26%				
Estimates						0.20	0.10	0.09	4.27
No adjustment for dividend taxes									
Transfer rates	23%	4%	1%	0%	0%				
Tax rates	16%	19%	22%	24%	30%				
Estimates						0.25	0.08	0.09	4.22
Medicaid and SSI in transfers									
Transfer rates	50%	11%	3%	1%	0%				
Tax rates	16%	19%	22%	24%	29%				
Estimates						0.25	0.08	0.18	3.44

Notes: The table summarizes tax-and-transfer rates by income quintile for different definitions of taxes, transfers, and income, as well as the corresponding estimates for the fiscal parameters.

B.1.3 Hours

We measure hours dispersion using the same sample as for the tax function estimation. We compute hours at the household level as weeks worked in the year (wkswork1) times usual hours per week (uhrsworkly), summed up over household members. We then regress household hours on: household type (head is single or head has married/unmarried partner living in the household), number of household members in different age groups, age, college education, race, and sex of the head. We only consider singles who work between 260 hours (Heathcote, Perri, and Violante 2010) and 4160 hours (that is, 80 hours for 52 weeks), and couples where total household hours are above 520 and each worker provides less than 4160 hours. We compute the variance of the residuals as our measure of hours dispersion. We proceed in the same way for wages (total labor income divided by total hours) and correlate the residuals.

B.2 Survey of Consumer Finances

We use the 2013 Survey of Consumer Finances (SCF) to compute income and wealth distributions. The SCF reports measures of wealth and oversamples the rich, offering a better description of the right tail of the income distribution than the CPS. The unit of observation is a family, defined as the economically dominant single person or couple (whether married or living together as partners) and all other persons in the household who are financially interdependent with that economically dominant person or couple. We restrict the sample to households where the head is between 25 and 60 and the household labor income and total income (as defined below) are above $5,000. Weights are used throughout.

Consistent with our definition in the CPS, we define labor income as the sum of wage and salary income (X5702), income from sole proprietorship and farm (X5704), income from other businesses or investments, net rent, trusts, or royalties (X5714), and income from unemployment or worker's compensation (X5716). Total income consists of labor income, income from non-taxable investments such as municipal bonds (X5706), income from other interest (X5708), income from dividends (X5710), income from capital gains (X5712), and income from other sources (X5724). Net worth (networth) consists of all real and financial assets net of all debts.

B.3 Panel Study of Income Dynamics

We resort to the Panel Study of Income Dynamics (PSID) to compute moments of the earnings growth distribution. We focus on pre-tax labor earnings at the household level. We use surveys from 1970 to 1992, as income definitions change after 1992 and the survey frequency becomes biannual after 1997. To be consistent with our data treatment in CPS and SCF, we include the labor incomes of head and wife, consisting of: the labor part of farm income and business income, wages, bonuses, overtime, commissions, professional practice, plus the labor part of income from roomers and boarders.[37] We further add: unemployment benefits received by the head; and the shares of business and farm income allocated to asset income. Before 1976 the asset part of

[37] "Head" and "wife" are the PSID terminology for this time period.

Table B.5: Earnings growth distribution: robustness

Moments	Stand. deviation	P90-10	Skewness	Kurtosis
Benchmark	0.33	0.60	-0.37	12.48
Drop pre-1976	0.33	0.60	-0.32	12.23
No asset part	0.33	0.61	-0.37	11.33
No UI	0.34	0.62	-0.37	12.24
Not control. edu.	0.33	0.60	-0.36	12.71
Cutoff at $1,500	0.37	0.62	-0.25	16.27
Cutoff at $10,000	0.31	0.58	-0.23	11.52

Notes: This table summarizes moments of the earnings growth distribution for different earnings definitions.

business and farm income is bracketed and we assign the mid point of the interval. We translate incomes into real values deflating by the CPI. Income growth is computed as log differences. We use households where the head is between ages 25 and 60 and the household labor income is above $5,000 in 2012 dollars. We use only the representative Survey Research Center sample.

We compute moments of the earnings growth distribution while keeping the household composition fixed; that is, requiring the head and wife to be the same individuals for three consecutive years. We use residualized earnings; that is, we regress earnings on year, age of head, and college education of head, and compute growth rates of the residuals. Table B.5 presents the results, as well as several robustness checks where: (1) we drop observations before 1976 for whom parts of income is bracketed: (2) we exclude the asset parts of business and farm income; (3) we exclude unemployment benefits; (4) we do not control for education; and (5) we use two other income thresholds typically used in the literature: $1,500, as in De Nardi, Fella, and Paz-Pardo (2019), and $10,000, as in Heathcote and Tsujiyama (2021).

B.4 National Income and Product Accounts

We use data from the National Income and Product Accounts (NIPA) to calibrate consumption taxes, capital taxes, and government spending.[38]

Following Bhandari and McGrattan (2021) and Bhandari, McGrattan, and Yao (2020), we estimate consumption taxes as the ratio of revenues from sales and excise taxes (NIPA table 3.2 lines 5, 6, 7; table 3.3 lines 7, 8, 10), raised at the federal, state, and local level, over total personal consumption expenditure (NIPA table 2.1 line 29). This yields a consumption tax of 6.26% in 2012.

We calibrate capital taxes τ_k to match the ratio of revenues from capital taxes to GDP. Our measure of revenues adds revenues from the corporate income tax (NIPA table 3.1 line 5) to an estimate for revenues from dividend taxes, which is obtained by multiplying income from dividends (NIPA table 2.1 line 15) with a tax rate of 13.3%, as estimated in Bhandari and

[38]NIPA data is from the U.S. Bureau of Economic Analysis, accessed on November 6, 2022 with the last data revision from October 27, 2022.

McGrattan (2021) and Bhandari, McGrattan, and Yao (2020).[39] Note that we account for business income in labor income, which is taxed at personal income tax rates. Also, a large portion of interest income accrues in tax deferred accounts (Bhandari et al. 2021), so we assume zero revenues from interest income. GDP is given in NIPA in Table 1.1.5 line 1. We obtain a ratio of revenues from capital taxes to GDP of 2.69%, resulting in an estimated capital tax rate of $\tau_k = 29.84\%$.

We measure government revenues as total receipts (NIPA table 3.1 line 1) net of income receipts on assets, current transfer receipts, and current surplus of government enterprises (NIPA table 3.1 lines 10, 15, 19), for which there is no counterpart in the model. Government revenue-to-GDP ratio amounts to 24.3% in the data, a number well matched in our model. Note that total expenditure (NIPA table 3.1 line 20) amounts to 34.9% of GDP, the difference between expenditure and revenues being accounted for by deficits. As the calibration is in steady state, we choose to target total revenues rather than total expenditure.[40]

Government revenues are split between interest payments, transfers, and public spending in the model. Interest payments (NIPA table 3.1 line 27) account for 4.33% of GDP, versus 2.0% in the model. Interest payments in NIPA include the service of the debt at the state and city levels, which is typically charged with higher interest rates. We compute transfers at the aggregate level as the sum of food stamps (NIPA table 3.12 line 21), refundable tax credits (NIPA table 3.12 line 25), family assistance (NIPA table 3.12 line 35), general assistance (NIPA table 3.12 line 37), energy assistance (NIPA table 3.12 line 38), and other assistance (NIPA table 3.12 line 39). This computation results in a transfer-to-output ratio of 1.42%—slightly larger than in the model, which is to be expected as the model is calibrated to working-age households with positive labor income.

B.5 Flow of Funds

The target for the debt-to-GDP ratio is computed as follows. From Table D3 in Financial Accounts of the United States, we sum federal government, state and local debt. Dividing by GDP yields a debt-to-GDP ratio of 98.6%.[41]

C Quantitative model: Calibration and solution

C.1 Tax function

Figure C.1 compares average and marginal tax rates implied by the log-linear tax function and the new tax function. In the new tax function, we set $\theta = 0.08$ and $\lambda = 0.25$, as in the

[39] Their estimate is for 2007, whereas we calibrate to the 2012 U.S. economy. Statutory rates for both ordinary and qualified dividends were similar in these years.

[40] An equivalent interpretation of the model is that the government receives, for free, an extra 10% of GDP in public goods, which are produced by the rest of the world. Whether that inflow is maintained or not after the tax reform is irrelevant to the optimal tax reform, as G is not valued.

[41] The Financial Accounts are released by the Federal Reserve Board. We accessed the data on November 7, 2022 with the data having been released on September 9, 2022.

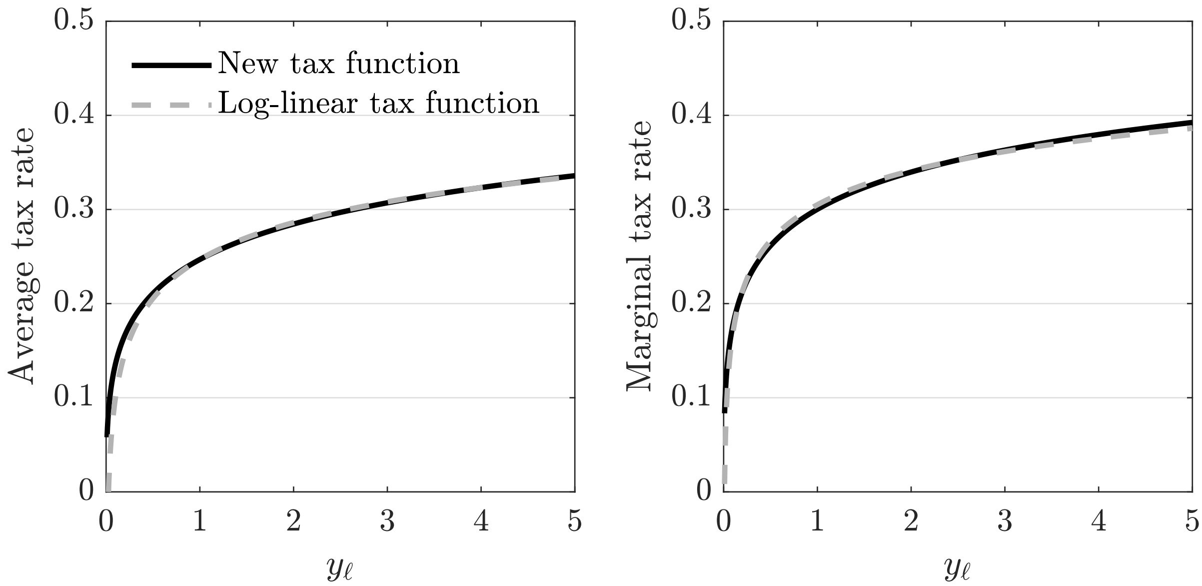

Figure C.1: New income tax function: comparison to log-linear function

Notes: This figure compares average and marginal tax rates implied by the log-linear tax function and the new tax function for $\tau = \theta = 0.08$. Labor income is plotted relative to mean income.

calibration. The corresponding parameters in the log-linear tax function are $\tau = 0.08$ and $\lambda = 1 - 0.25$.

C.2 Numerical solution

C.2.1 Steady state

To solve for the steady state of the economy, we need to find the real interest rate r that clears the capital market and the level of transfers m that clears the government budget constraint. We explain next how we do this.

0. Set grids for assets $\vec{a}$ and productivity levels $\vec{z}$. Let $N_a = 300$ and $N_z = 19$ be the number of points in each grid, respectively. Compute the transition matrix of productivities $\pi_z(z', z)$ using Farmer and Toda (2017).[42]

1. Guess values for the interest rate r and the transfer parameter m. From the firm's first order conditions, we can compute the wage w implied by the guessed r.

2. Solve for household policies by value function iteration. In particular, for a given guess, guess a value function $V(a, z)$ and update the value function as

$$\hat{V}(a, z) = \max_{a' \geq \underline{a}, n \in [0,1]} \left\{ \frac{c^{1-\sigma}}{1-\sigma} - B\frac{n^{1+\varphi}}{1+\varphi} + \beta \sum_{z' \in \vec{z}} \pi_z(z', z) V(a', z') \right\}$$
$$\text{s.t. } (1+\tau_c)c + a' \leq wzn + (1+r)a - \mathcal{T}(wzn, ra).$$

[42]We use the Matlab package provided at https://alexisakira.github.io/discretization/.

Iterate until $\left\|\hat{V} - V\right\| < \varepsilon^V$. We use $\varepsilon^V = 1e - 10$.

3. Compute the stationary measure implied by the optimal policies of step 2. In particular, for a given guess $\mu(a,z)$, compute implied measure $\hat{\mu}(a,z)$ as

$$\hat{\mu}(a_{i'}, z_{j'}) = \sum_{i=1}^{N_a} \sum_{j=1}^{N_z} \mathbb{L}\left\{a_{i'} = a'(a_i, z_j)\right\} \pi_z(z_{j'}, z_j) \mu(a_i, z_j)$$

   where $\mathbb{L}$ computes a linear interpolation: $\mathbb{L}(a_i, a') = \mathbb{I}\left(a' \in (a_{i-1}, a_i]\right) \frac{a - a_{i-1}}{a_i - a_{i-1}}$. Iterate until $\|\hat{\mu} - \mu\| < \varepsilon^\mu$. We use $\varepsilon^\mu = 1e - 11$.

4. Compute asset market clearing error: $E^K = A - D - \hat{K}$ where $A = \int a d\mu(a,z)$ is households' asset holdings and $\hat{K} = L\left(\frac{r+\delta}{1-\alpha}\right)^{-1/\alpha}$ is capital demand given interest rate and labor supply $L = \int n(a,z) d\mu(a,z)$. Also compute government budget constraint error $E^G = G + rD - \int \mathcal{T}\left(wzn(a,z), ra\right) d\mu(a,z) + \tau_c \int c(a,z) d\mu(a,z)$. Let $\mathcal{E}(X) \equiv \left(E^K, E^G\right)$ collect the two errors given the guess $X \equiv \{m, r\}$. An equilibrium can be written as

$$\mathcal{E}(X) = 0. \tag{27}$$

   We solve for X in equation (27) using a quasi-Newton method.

C.2.2 One transition

We assume a once-and-for-all fiscal reform, where the two tax parameters, λ and θ ,and the transfer phase-out, ξ, jump to their new values. The transfer level $\{m_t\}$ adjusts every period to clear the government budget constraint. We assume that the economy has converged to its new steady state $\bar{T}$ periods after the shock. We first compute the new steady state as described in Section C.2.1 and obtain the value function $\bar{V}(a,z)$ and the equilibrium vector $\bar{X} = (\bar{m}, \bar{r})$ of the new steady state. As the economy has converged in $\bar{T}$, we know that the value function at $t = \bar{T}$ equals its steady-state value $V_{\bar{T}}(a,z) = \bar{V}(a,z)$. We also know that the measure at time $t = 1$ is equal to the initial steady-state value $\mu_1(a,z) = \mu(a,z)$. Then, given a guess for transfers and interest rates $\{m_t, r_t\}_{t=1}^{\bar{T}}$ such that $(m_{\bar{T}}, r_{\bar{T}}) = (\bar{m}, \bar{r})$, we solve the household problem backwards, the measure μ_t forward, and iterate on the sequence $\{m_t, r_t\}_{t=1}^{\bar{T}}$ using a quasi-Newton algorithm to clear markets.

C.2.3 Global optimum

To find the optimal tax reform, we rely on the TikTak Global Optimization multistart algorithm—see Arnoud, Guvenen, and Kleineberg (2022) for a detailed description.[43]

We optimize on the triplet $\{\theta, \lambda, \xi\}$. For each triplet, we compute the transition as described in Section C.2.2. The algorithm looks for the triplet that maximizes welfare at the

[43] We use the Fortran package provided at https://github.com/serdarozkan/TikTak.

Table C.1: Income risk in the model

Moments	Stand. deviation	P90-10	Skewness	Kurtosis
Data	0.33	0.60	-0.37	12.48
Model	0.34	0.62	-0.31	13.25

Notes: Income risk in the data (PSID) and in the model.

Table C.2: Distribution of hours in the model

log hours	Q1	Q2	Q3	Q4	Q5	$v(\log h)$
Data	-0.47	-0.03	0.05	0.13	0.33	0.097
Model	-0.41	-0.06	0.07	0.16	0.24	0.063

Notes: Average log hours per quintile minus mean of log hours, sorting households per hours worked, variance of log hours. Empirical moments are computed using CPS date (see Section B.1).

implementation of the fiscal reform; that is,

$$\int V_1(a, z) d\mu(a, z).$$

We start local optimizers using the best 40 points in a Sobol sequence of 672 points. The local optimizer is Nelder Mead.

C.3 Income risk

Table C.1 compares the income risk moments in the model to their targeted data counterparts.

C.4 Hours

Table C.2 presents two measures of dispersion: average log hours per quintile of hours worked relative to the average log hours across all households, and variance of log hours. The correlation of hours worked with hourly wages is close to zero: it is slightly positive at 0.04 in the model, vs. slightly negative at -0.05 in the data.

C.5 Model validation: Wealth effects and elasticities

Wealth effects.—Golosov et al. (2021) combine data covering the universe of U.S. taxpayers from 1999 to 2016 with data on winnings in state lotteries. There are 45 U.S. states that run a state lottery and winnings above \$600 are reported to tax authorities, as lottery winnings are considered taxable income. They restrict the sample to wins of more than \$30K. They estimate wealth effects on labor earnings during the five years following a lottery win, and report average effects.

We replicate this estimation strategy in the model as follows. Starting from the steady state distribution, we increase each household's assets by the size of the lottery win. We perform our exercise in partial equilibrium, where prices and taxes are held constant–as the mass of lottery winners is negligible in the data and winnings are not financed from taxes. We compute average labor earnings of households for the next five years using non-stochastic simulation, similar to Young (2010). As control group, we also simulate a panel of households not experiencing a wealth shock. The difference between the labor earnings of the two groups gives the change in labor earnings due to the wealth shock.

Section 3.5 reports estimates from Golosov et al. (2021) as well as the model implied responses for two win sizes. The first size is the average lottery win size of \$180,000 in 2015 U.S. dollars. Large wins are defined as being above \$1,000,000, but their distribution is not provided. We use a large win size of \$2,000,000, which roughly corresponds to the present discounted value of a (annual) perpetuity of \$28,200, discounted with the calibrated after-tax interest rate of 1.4%.

Labor elasticity at the top.—We follow Kindermann and Krueger (2022) to evaluate labor elasticities in the model. Starting from the calibrated steady state, we assume an unexpected 1%-increase of labor tax rates for all households. We consider different cases varying the persistence of the tax change: from a one-period change to a permanent change. The experiment is conducted in partial equilibrium, and without adjusting the government's budget constraint. We report the labor elasticity of the top-1% income group, using their labor response at the moment of the shock.

D Quantitative model: Results

D.1 Optimal steady state

The fiscal plan that optimizes steady-state welfare is: $m = 0.21$, $\theta = 0.11$, $\lambda = 0.28$, and $\xi = 5.63$. In line with Bakış, Kaymak, and Poschke (2015), the optimal plan is more generous when incorporating transitions. Transfers at the bottom are only \$17,500, compared to \$19,800 when including transitions.

D.2 Consumption equivalents and welfare decomposition

D.2.1 Consumption equivalents

Consumption equivalents $\gamma(a,z)$ are computed as the increase in consumption in the status quo which would make household (a,z) indifferent between the status quo and the tax reform. More formally, one can write $V(a,z)$, the life-time utility of a household with state (a,z) in the status quo, as

$$V(a,z) = \mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t \left\{ \frac{c_t^{1-\sigma}}{1-\sigma} - B \frac{n_t^{1+\varphi}}{1+\varphi} \right\} |a,z \right]$$

Table D.1: Welfare decomposition

Source of welfare gains	Total	Consumption	Leisure
Efficiency	4.6%	-115.5%	120.0%
Redistribution	17.9%	14.6%	3.3%
Insurance	77.5%	84.5%	-7.0%

Notes: This table decomposes the welfare gains in three components: efficiency, redistribution, and insurance. Each component is further decomposed in terms of consumption and leisure.

where expectation $\mathbb{E}_0$ is taken over future paths of shocks given individual states at the moment of the tax reform.

For a given policy $\boldsymbol{\tau}$, let $V_1^{\boldsymbol{\tau}}(a,z)$ be the utility of household (a,z) when implementing the reform. We compute $\gamma(a,z)$ such that

$$\mathbb{E}_0\left[\sum_{t=0}^{\infty}\beta^t\left\{\frac{[(1+\gamma(a,z))c_t]^{1-\sigma}}{1-\sigma}-B\frac{n_t^{1+\varphi}}{1+\varphi}\right\}|a,z\right]=V_1^{\boldsymbol{\tau}}(a,z).$$

Consumption equivalents are thus equal to

$$\gamma(a,z)=-1+\left[1+\frac{V_1^{\boldsymbol{\tau}}(a,z)-V(a,z)}{\mathcal{C}(a,z)}\right]^{\frac{1}{1-\sigma}}, \tag{28}$$

where $\mathcal{C}(a,z)=\mathbb{E}_0\left[\sum_{t=0}^{\infty}\beta^t\left\{\frac{c_t^{1-\sigma}}{1-\sigma}\right\}|a,z\right]$ can be computed using a simple iteration of policy functions. Finally, we aggregate over the distribution of consumption equivalents using the measure in the status quo.

D.2.2 Welfare decomposition

We follow the new decomposition proposed by Bhandari et al. (2022).[44] Table D.1 reports the contribution of aggregate efficiency, redistribution, and insurance to the total welfare gains generated by the tax reform described in Section 4. Each component is further decomposed in terms of leisure and consumption.

D.3 Optimal log-linear plan

We compute the optimal log-linear tax plan using a labor tax function as in the analytical section: $\mathcal{T}(y)=y-\lambda(y/\bar{y})^{1-\tau}$. The economy is calibrated in steady state as described in Section 3, and a planner implements a one-time fiscal reform where the benchmark taxes and transfers are replaced by a log-linear function with constant income-tax progressivity τ. As in Section 4, we take into account transitions for welfare computations. The optimal plan features $\tau=0.39$. Table 5 reports average and marginal tax rates across income quintiles.

[44] We thank the authors for sharing detailed notes on how to implement this decomposition in a related framework.

Table D.2: Optimal $t\&T$ plans with phasing-out under alternative calibrations

	m	θ	ξ	λ
Income distribution				
Benchmark	0.23	0.14	3.41	0.36
No Pareto	0.22	0.06	2.83	0.37
Richer Poor	0.14	0.15	1.97	0.34
Income risk				
No Pareto Normal	0.22	0.06	2.46	0.39
No Pareto Persistent	0.23	0.05	2.34	0.41
Fiscal Space				
Low Spending	0.25	0.16	3.18	0.32
Preferences				
Low σ & φ	0.20	0.13	4.04	0.31

Notes: This table presents statistics of the optimal $t\&T$ plan with phasing-out for the benchmark economy and 7 alternative calibrations.

D.4 Robustness

Table D.2 presents the parameters of the optimal $t\&T$ system with phasing-out, for the seven alternative calibrations described in Section 5. We also report the benchmark for comparison. Transfers are rescaled to be comparable to the benchmark economy.

The main take-aways are in line with those for the optimal UBI presented in Table 7. The right tail of the income distribution primarily alters progressivity θ, while the left tail changes transfers m. Lower spending increases both optimal m and θ, while higher Frisch elasticity/lower risk-aversion tends to lower both m and θ.

E Data Availability

Code replicating the tables and figures in this article can be found in Ferriere et al. (2023) in the Harvard Dataverse, https://doi.org/10.7910/DVN/ZQVFEZ.

References

Acikgöz, Ömer, Marcus Hagedorn, Hans Holter, and Yikai Wang (2022). "The Optimum Quantity of Capital and Debt." *Working Paper*.

Aiyagari, S. Rao (1994). "Uninsured idiosyncratic risk and aggregate saving." *The Quarterly Journal of Economics* 109.3, pp. 659–684.

Aoki, Shuhei and Makoto Nirei (2017). "Zipf's Law, Pareto's Law, and the Evolution of Top Incomes in the United States." *American Economic Journal: Macroeconomics* 9.3, pp. 36–71.

Arnoud, Antoine, Fatih Guvenen, and Tatjana Kleineberg (2022). "Benchmarking Global Optimizers." *Working Paper.*

Attanasio, Orazio and Luigi Pistaferri (2014). "Consumption inequality over the last half century: some evidence using the new PSID consumption measure." *The American Economic Review* 104.5, pp. 122–26.

Ayaz, Mehmet, Lea Fricke, Clemens Fuest, and Dominik Sachs (2022). "Who should bear the burden of COVID-19 related fiscal pressure? An optimal income taxation perspective." *Working Paper.*

Badel, Alejandro, Mark Huggett, and Wenlan Luo (2020). "Taxing top earners: a human capital perspective." *The Economic Journal* 130.629, pp. 1200–1225.

Bakış, Ozan, Barış Kaymak, and Markus Poschke (2015). "Transitional dynamics and the optimal progressivity of income redistribution." *Review of Economic Dynamics* 18.3, pp. 679–693.

Benabou, Roland (2000). "Unequal societies: Income distribution and the social contract." *The American Economic Review* 90.1, pp. 96–129.

Bhandari, Anmol and Ellen R. McGrattan (2021). "Sweat equity in US private business." *The Quarterly Journal of Economics* 136.2, pp. 727–781.

Bhandari, Anmol, Ellen R. McGrattan, and Yuki Yao (2020). "Taxation of Paid-and Self-Employment." *Working Paper.*

Bhandari, Anmol, David Evans, Mikhail Golosov, and Thomas J. Sargent (2021). "Inequality, Business Cycles, and Monetary-Fiscal Policy." *Econometrica* 89.6, pp. 2559–2599.

— (2022). "Efficiency, Insurance, and Redistribution Effects of Government Policies." *Working Paper.*

Boar, Corina and Virgiliu Midrigan (2022). "Efficient redistribution." *Journal of Monetary Economics* 131, pp. 78–91.

Boerma, Job and Ellen R. McGrattan (2020). "Quantifying Efficient Tax Reform." *Working Paper.*

Brüggemann, Bettina (2021). "Higher taxes at the top: The role of entrepreneurs." *American Economic Journal: Macroeconomics* 13.3, pp. 1–36.

Carroll, Daniel, André Victor Doherty Luduvice, and Eric R. Young (2022). "Optimal Fiscal Reform with Many Taxes." *Working Paper.*

Chang, Yongsung and Yena Park (2021). "Optimal taxation with private insurance." *The Review of Economic Studies* 88.6, pp. 2766–2798.

Chetty, Raj (2012). "Bounds on elasticities with optimization frictions: A synthesis of micro and macro evidence on labor supply." *Econometrica* 80.3, pp. 969–1018.

Colombino, Ugo and Nizamul Islam (2022). "Combining Microsimulation and Numerical Maximization to Identify Optimal Tax-Transfer Rules." *IJM* 15.2, pp. 4–43.

Conesa, Juan Carlos and Dirk Krueger (2006). "On the optimal progressivity of the income tax code." *Journal of Monetary Economics* 53.7, pp. 1425–1450.

Conesa, Juan Carlos, Bo Li, and Qian Li (2021). "A Quantitative Evaluation of Universal Basic Income." *Working Paper.*

Daruich, Diego and Raquel Fernández (2022). "Universal Basic Income: A Dynamic Assessment." *Working Paper.*

De Nardi, Mariacristina, Giulio Fella, and Gonzalo Paz-Pardo (2019). "Nonlinear household earnings dynamics, self-insurance, and welfare." *Journal of the European Economic Association.*

Diamond, Peter A. (1998). "Optimal income taxation: an example with a U-shaped pattern of optimal marginal tax rates." *The American Economic Review*, pp. 83–95.

Dyrda, Sebastian and Marcelo Pedroni (2022). "Optimal fiscal policy in a model with uninsurable idiosyncratic shocks." *Forthcoming in The Review of Economic Studies.*

Farhi, Emmanuel and Iván Werning (2013). "Insurance and taxation over the life cycle." *The Review of Economic Studies* 80.2, pp. 596–635.

Farmer, Leland E. and Alexis Akira Toda (2017). "Discretizing nonlinear, non-Gaussian Markov processes with exact conditional moments." *Quantitative Economics* 8.2, pp. 651–683.

Feldstein, Martin S. (1969). "The effects of taxation on risk taking." *Journal of Political Economy* 77.5, pp. 755–764.

Ferriere, Axelle, Philipp Grübener, Gaston Navarro, and Oliko Vardishvili (2023). *Replication Data for: "On the Optimal Design of Transfers and Income-Tax Progressivity".* Version V1.

Findeisen, Sebastian and Dominik Sachs (2017). "Redistribution and insurance with simple tax instruments." *Journal of Public Economics* 146, pp. 58–78.

Flood, Sarah, Miriam King, Renae Rodgers, Steven Ruggles, J. Robert Warren, and Michael Westberry (2021). "Integrated Public Use Microdata Series, Current Population Survey: Version 9.0 [dataset]. Minneapolis, MN: IPUMS, 2021." *Dataset*.

Golosov, Mikhail, Maxim Troshkin, and Aleh Tsyvinski (2016). "Redistribution and social insurance." *The American Economic Review* 106.2, pp. 359–86.

Golosov, Mikhail, Michael Graber, Magne Mogstad, and David Novgorodsky (2021). "How Americans respond to idiosyncratic and exogenous changes in household wealth and unearned income." *Working Paper*.

Greenwood, Jeremy, Zvi Hercowitz, and Gregory W. Huffman (1988). "Investment, capacity utilization, and the real business cycle." *The American Economic Review*, pp. 402–417.

Guner, Nezih, Remzi Kaygusuz, and Gustavo Ventura (2021). "Rethinking the Welfare State." *Working Paper*.

Guner, Nezih, Martin Lopez-Daneri, and Gustavo Ventura (2016). "Heterogeneity and Government revenues: Higher taxes at the top?" *Journal of Monetary Economics* 80, pp. 69–85.

Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song (2021). "What Do Data on Millions of US Workers Reveal about Life-Cycle Earnings Risk?" *Econometrica* 89.5, pp. 2303–2339.

Habib, Bilal (2018). "How CBO Adjusts for Survey Underreporting of Transfer Income in Its Distributional Analyses." *Working Paper*.

Heathcote, Jonathan, Fabrizio Perri, and Giovanni L. Violante (2010). "Unequal we stand: An empirical analysis of economic inequality in the United States, 1967–2006." *Review of Economic Dynamics* 13.1, pp. 15–51.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante (2014). "Consumption and labor supply with partial insurance: An analytical framework." *The American Economic Review* 104.7, pp. 2075–2126.

— (2017). "Optimal tax progressivity: An analytical framework." *The Quarterly Journal of Economics* 132.4, pp. 1693–1754.

Heathcote, Jonathan and Hitoshi Tsujiyama (2021). "Optimal income taxation: Mirrlees meets Ramsey." *Journal of Political Economy* 129.11, pp. 3141–3184.

Holter, Hans A., Dirk Krueger, and Serhiy Stepanchuk (2019). "How do tax progressivity and household heterogeneity affect Laffer curves?" *Quantitative Economics* 10.4, pp. 1317–1356.

Hubmer, Joachim, Per Krusell, and Anthony A. Smith (2020). "Sources of US wealth inequality: Past, present, and future." *NBER Macroeconomics Annual 2020, Volume 35*. University of Chicago Press.

Jaimovich, Nir, Itay Saporta-Eksten, Ofer Setty, and Yaniv Yedid-Levi (2022). "Universal basic income: Inspecting the mechanisms." *Working Paper*.

Kapička, Marek (2013). "Efficient allocations in dynamic private information economies with persistent shocks: A first-order approach." *Review of Economic Studies* 80.3, pp. 1027–1054.

Kindermann, Fabian and Dirk Krueger (2022). "High Marginal Tax Rates on the Top 1 Percent? Lessons from a Life-Cycle Model with Idiosyncratic Income Risk." *American Economic Journal: Macroeconomics* 14.2, pp. 319–66.

Krueger, Dirk and Alexander Ludwig (2016). "On the optimal provision of social insurance: Progressive taxation versus education subsidies in general equilibrium." *Journal of Monetary Economics* 77, pp. 72–98.

Lin, Daniel (2022). "Methods and Assumptions of the CPS ASEC Tax Model." *Working Paper*.

Lopez-Daneri, Martin (2016). "NIT picking: The macroeconomic effects of a Negative Income Tax." *Journal of Economic Dynamics and Control* 68, pp. 1–16.

Luduvice, André Victor Doherty (2021). "The Macroeconomic Effects of Universal Basic Income Programs." *Working Paper*.

Mankiw, N. Gregory, Matthew Weinzierl, and Danny Yagan (2009). "Optimal taxation in theory and practice." *Journal of Economic Perspectives* 23.4, pp. 147–74.

Meyer, Bruce D., Derek Wu, Grace Finley, Patrick Langetieg, Carla Medalia, Mark Payne, and Alan Plumley (2020). "The Accuracy of Tax Imputations: Estimating Tax Liabilities and Credits Using Linked Survey and Administrative Data." *Working Paper*.

Mirrlees, James A. (1971). "An exploration in the theory of optimum income taxation." *The Review of Economic Studies* 38.2, pp. 175–208.

Panel Study of Income Dynamics (2021). "Produced and distributed by the Survey Research Center, Institute for Social Research, University of Michigan, Ann Arbor, MI (2021)." *Dataset*.

Park, Yena (2022). "Tax Incidence and Optimal Taxation with General Equilibrium Effects and Transition." *Working Paper*.

Peterman, William B. (2016). "The effect of endogenous human capital accumulation on optimal taxation." *Review of Economic Dynamics* 21, pp. 46–71.

Rauh, Christopher and Marcelo Rodrigues Santos (2022). “How Do Transfers and Universal Basic Income Impact the Labor Market and Inequality?” *Working Paper*.

Rauh, Joshua and Ryan J. Shyu (2022). “Behavioral responses to state income taxation of high earners: evidence from California.” *Working Paper*.

Saez, Emmanuel (2001). “Using elasticities to derive optimal income tax rates.” *The Review of Economic Studies* 68.1, pp. 205–229.

Stantcheva, Stefanie (2017). “Optimal taxation and human capital policies over the life cycle.” *Journal of Political Economy* 125.6, pp. 1931–1990.

Toda, Alexis Akira and Kieran Walsh (2015). “The double power law in consumption and implications for testing Euler equations.” *Journal of Political Economy* 123.5, pp. 1177–1200.

Young, Eric R. (2010). “Solving the incomplete markets model with aggregate uncertainty using the Krusell–Smith algorithm and non-stochastic simulations.” *Journal of Economic Dynamics and Control* 34.1, pp. 36–41.