STUDY CENTER GERZENSEE

On the Optimal Design of Consumption Taxes

Michael Barczay

Working Paper 25.03

On the optimal design of consumption taxes

Michael Barczay

European University Institute & Study Center Gerzensee

August 21, 2025

Abstract

This paper studies the optimal design of differentiated consumption taxes in the presence of progressive labor income taxes and capital income taxation. A quantitative heterogeneous-agent model with non-homothetic preferences and uninsurable idiosyncratic risk is estimated using US consumption and price data to match expenditure patterns across the income distribution. Solving the Ramsey problem in which the government jointly chooses labor income and commodity taxes, the optimal policy prescribes a subsidy on necessities of -52% and a positive tax of 7% on luxuries, accompanied by a reduction in labor tax progressivity. Three mechanisms account for these results: subsidized necessities provide consumption insurance, taxation of luxuries acts as an implicit tax on existing wealth, and differentiated rates strengthen labor supply incentives among highly productive households.

Keywords: Heterogeneous Agents, Fiscal Policy, Optimal Taxation, Redistribution
JEL Classification: E21, E62, H21, H23, H31

I am indebted to my advisors Alexander Monge-Naranjo and Russell Cooper for their advice and support. I would also like to thank Árpád Ábrahám, Axelle Ferriere, Lukas Nord, Ernesto Pasten, Sofía Sierra Vásquez, Joel Slemrod, Kjetil Storesletten, Vytautas Valaitis, and workshop participants at the EUI, the SEA End-of-Year Conference 2022, Banco Central de Chile, SAEe 2023, Study Center Gerzensee, MMF Annual Conference 2024, and the Doctoral Workshop on Quantitative Dynamic Economics (AMSE) 2024 for comments and suggestions that helped to greatly improve this work. Any remaining errors are my own.
Correspondence: Michael Barczay, European University Institute, Department of Economics, Via delle Fontanelle 18, 50014 San Domenico di Fiesole, Italy. Email: michael.barczay@eui.eu.

1 INTRODUCTION

The rise in income and wealth inequality in recent decades has sparked extensive debate about possible policies to counteract these trends. Tax policy has emerged as a particularly potent tool to do so. Numerous tax reform proposals have been put forth, ranging from higher tax rates for the top 1% earners (Kindermann and Krueger, 2022) to higher transfers (Ferriere et al., 2023), or progressive wealth taxes (Boar and Midrigan, 2022). In this paper, I study consumption taxes as an alternative instrument and explore their optimal design when different goods and services can be subject to differentiated tax rates, while considering the government's availability of other tax instruments.

To do so, I develop a quantitative model featuring heterogeneous households with non-homothetic preferences over 11 categories of consumption goods and services, ranging from groceries and utilities to food consumed in restaurants and recreational expenses. I calibrate the model to US data by estimating the numerous parameters governing the behavior of agents in my economy using US household data and show that my model matches the heterogeneous consumption behavior across the 11 categories. To solve for the optimal tax design, I adopt a quantitative Ramsey approach, allowing the government to set taxes on different consumption goods and services as well as on concurrent parametrized tax functions. Crucially, the government takes into account the full transition from the initial steady state to the final steady state when computing the welfare effects of tax reforms.

I derive two main results from this model. First, it is welfare maximizing to provide a generous subsidy for the consumption of necessities and to impose a positive tax rate on luxuries. This policy reform reduces consumption inequality, acts as a tax on the initial wealth of very wealthy households, and incentivizes a relatively higher labor supply of very productive households. Second, I show that by allowing the planner to jointly optimize consumption and labor taxes, the progressivity of labor taxes is reduced drastically from the *status quo*. Increasing the differential between the rate on necessities and luxuries provides consumption insurance and can therefore be compensated by reducing the progressivity of the labor income tax. Crucially, while the labor income tax base is solely labor income, the consumption tax can implicitly target high-wealth households, which explains why the two instruments are imperfect substitutes. More generally, the reform can be understood as a shift from providing labor income insurance to a policy that directly insures consumption. Finally, I show that such a shift of redistributing using the labor income tax to increasing the rate differential between necessities and luxuries generates efficiency gains by incentivizing highly productive households to increase their labor supply.

The optimal design of consumption taxes is a problem along two dimensions. The first dimension concerns the fundamental question of the desirability of consumption taxes. Opponents of consumption taxes argue that low-income households, with their lower savings rates, bear a disproportionate burden when subject to a positive consumption tax. Proponents of consumption taxes contend that they represent and implicit lump-sum tax with minimal distortions on existing wealth

at the time of the implementation (Nishiyama and Smetters, 2005). The second dimension involves determining the optimal differentiation of tax rates among various goods and services. Advocates for differentiated consumption taxes cite inequality and fairness arguments. If, for instance, luxury goods are taxed at a higher rate than necessity goods, then this can reduce consumption inequality by implying higher average consumption tax rates for high-income households. In contrast, Ramsey (1927) or more recently Chari and Kehoe (1999) have argued that goods with low price or income elasticity should be applied a higher rate since this induces minimal distortions. Notably, these two predictions may conflict, given that low-income households predominantly consume inelastic necessities. Contrary to this, the public finance literature has often considered uniform consumption taxes to be optimal since differentiated rates can easily be replicated by optimal income taxes, without distorting consumption decisions (Atkinson and Stiglitz, 1976).
Finally, the optimal design of consumption taxes also hinges crucially on the existence and design of concurrent tax instruments. If alternative tax instruments are already designed to take into account redistributive goals, then the emphasis of consumption taxes may lay on efficiency considerations. Instead, if other tax instruments leave room for further redistribution, then consumption taxes may need to be designed such that goods and services primarily consumed by low-income households are taxed at a reduced rate.

In this paper, I propose a quantitative incomplete markets model to study the optimal level of consumption taxes, the optimal degree of rate differentiation and the interplay of consumption taxes with other tax instruments. I develop a general equilibrium model with uninsurable idiosyncratic risk and various consumption goods and services. The backbone of the model constitute heterogeneous households with non-homothetic constant elasticity of substitution (CES) preferences over a variety of consumption goods and services (Comin et al., 2021). Crucially, households consume not only pure consumption goods but also an outside good, which captures non-consumption expenditures, such as durables, housing, or insurance payments. While the analysis does not directly examine these expenditures, their inclusion accounts for the empirical pattern that not all expenditures are subject to consumption taxes and that the split between (taxed) consumption- and other (untaxed) expenditures is non-homothetic in itself. With these outside expenditures explicitly included in the model, differences in total consumption expenditures and with it differences in consumption baskets are induced by three layers of non-homothetic behavior: First, differences in savings rates induce differences in total expenditures. Second, the non-homothetic inclusion of the outside good generates differences in the split of expenditures across outside expenditures and consumption expenditures. Finally, the non-homothetic CES utility function generates empirically realistic differences in consumption baskets across the income distribution.
The model parameters are estimated using microdata from the Consumption Expenditure Survey covering the years 2001 to 2019. The non-homothetic CES preferences generate household demand equations for various consumption goods and services and using the Generalized Method of Moments, I recover the parameters of the non-homothetic CES. Calibrating the model with the estimated parameters, I show that the model captures well the consumption behavior of US house-

holds across the income distribution. In a second step, I estimate various parameters of the model to match data moments related to the aggregate US data moments and key numbers related to the US tax system.

Compared to previous work on the topic, my modeling and estimation approach offers several advantages. First, by generating realistic consumption behavior across various goods and services, while also incorporating household labor supply and savings decisions, the analysis enables a nuanced examination of the fundamental equity-efficiency trade-offs inherent in the optimal design of consumption taxes. Second, the endogenous categorization of goods and services into necessities and luxuries obviates the need for *a priori* classifications, as has often been the case in the literature (see, for instance, Conesa et al. (2020) or Parodi (2023)). Third, the adoption of a Ramsey approach to optimal taxation allows for a comprehensive evaluation of consumption tax reforms, taking into account the interplay between consumption taxes and alternative tax instruments. These alternative tax instruments are represented within a flexible parametric framework, which facilitates the evaluation of tax reforms by explicitly incorporating empirically grounded efficiency and equity considerations. Finally, the dynamic structure of my model allows for the analysis of tax reforms by taking into account the long-lived transition of the economy from one steady state to another – an aspect that can significantly shape the welfare conclusions (see e.g. Auerbach and Kotlikoff, 1987).
I employ the estimated model as a laboratory to solve for and analyze the optimal design of consumption taxes. The analysis proceeds in two steps. Initially, the concurrent tax instruments are held at their current *status quo*, and optimization is conducted solely with respect to the tax rates on necessities and luxuries. Subsequently, I allow the planner to jointly optimize the tax on various consumption goods and services as well as the schedule of the labor income tax.

Findings. Two main results emerge from solving the Ramsey Problem in which the government is allowed to optimize both consumption taxes and the design of the labor income tax. First, optimal consumption taxes are clearly non-uniform. Necessities, such as groceries, utilities or public transportation, are optimally taxed at -52%, effectively providing a subsidy for the consumption of these essential goods and services. Conversely, luxury goods and services, encompassing domestic services, restaurant food or recreational expenses, are optimally taxed at a rate of 7%.
Second, I show that these differentiated tax rates are welfare maximizing even under the optimal labor income tax. In fact, allowing the planner to jointly optimize consumption and labor income taxes yields optimal labor income taxes which are substantially less progressive. The large subsidy on necessity goods is, however, mirrored by an increase in the average labor income tax faced by households.
Several factors contribute to these findings. Taxing goods and services predominantly consumed by low-income households at a reduced rate enables a substantial reduction in consumption inequality: the consumption Gini coefficient is reduced by approximately 5 percentage points as compared to the calibrated initial steady state. These different consumption tax rates also provide

consumption insurance by subsidizing goods and services primarily consumed by households experiencing adverse idiosyncratic shocks to their productivity.

At the same time, the model generates households with large wealth holdings whose consumption basket is dominated by luxury goods and services, such as food consumed in restaurants or recreational expenses. Given the risk associated with labor income, these high-wealth households consume even more luxury goods as compared to households with identical labor income but less wealth. Consequently, levying a positive tax rate on luxury goods serves as an implicit tax on the considerable wealth positions of these high-wealth households. Since the government takes into account the transition when computing the welfare effects of policy reforms, the consumption tax provides a particularly potent tool to tax wealth with minimal distortions: At the time of the reform, high-wealth households cannot adjust their wealth holdings except by consuming and thereby paying the higher consumption tax on luxuries.

The fact that consumption insurance is implemented using the consumption tax allows for a less progressive labor income tax. As compared to providing insurance using the labor income tax, the consumption tax has the combined benefit of providing insurance while also introducing an implicit tax on wealth. While the non-uniform rates distort consumption decisions, the additional benefits of implicitly taxing wealth outweigh the cost induced by the consumption distortions. Additionally, I show that the policy reform positively affects labor supply by incentivizing highly productive households to increase their labor supply relative to unproductive households. On the one hand, the changes to the labor income tax reduce distortions induced by progressive labor income taxes. On the other hand, since under the optimal policy the goods and services favored by productive high-income households are taxed at 7%, these households are required to increase their labor supply in order to maintain their living standard. Taken together, these labor supply responses contribute to the efficiency gains of the policy reform. The combination of efficiency gains and increased insurance yield a welfare gain of 1.6% compared to the *status quo*, measured in terms of consumption equivalent variation. Furthermore, significantly more than 50% of the population benefit from this tax reform.

These results highlight the need to analyze consumption not merely has a homogeneous good but by considering the heterogeneity of consumption baskets across the income and wealth distribution. They also contribute to our understanding of the optimal degree of rate differentiation by suggesting that increasing the degree of rate differentiation can provide large welfare gains – even under optimally designed labor income taxes. Taken together, these findings highlight also highlight situations in which flat consumption taxes as in Atkinson and Stiglitz (1976) are dominated by heterogenous consumption tax rates in terms of aggregate welfare.

Related Literature. The analysis of tax policy using quantitative incomplete markets models has become a very active field and my work is directly related to it. Various strands of the literature have focused on different tax instruments. For instance, Conesa and Krueger (2006), Heathcote et al. (2017), Dyrda and Pedroni (2022), or Ferriere et al. (2023) provide an analysis of the optimal income tax progressivity. Starting with Aiyagari (1995) this methodological approach has also been used extensively to study capital income taxes. More recent work on capital income taxes includes the seminal work by Domeij and Heathcote (2004), or papers taking into account the complementarity of capital with certain types of labor (Kina et al., 2024; Bhattarai et al., 2022). Compared to these more focused papers, Boar and Midrigan (2022) provide a comprehensive analysis of various tax instruments but do not optimize over consumption taxes. I directly build on and contribute to this literature by focusing on consumption taxes as one of many tax instruments. While the mentioned papers often account for the existence of consumption taxes, they are usually not subject to the optimization problem of the planner. A limited number of papers have analyzed consumption taxes using a quantitative approach. Krusell et al. (1996) integrate a voting procedure into a neoclassical growth model with heterogenous housholds, to study the consumption and labor income taxes. One of the main findings is that when the population votes on only one tax instrument, then the income tax is preferred, whereas voters would prefer to have both income and consumption taxes in place if they get the opportunity to vote on multiple taxes. Macnamara et al. (2022) and Carroll et al. (2024) analyze the optimal design of consumption taxes in a life-cycle model with various tax instruments and find very positive consumption tax rates under the optimal policy. Compared to my work, their models feature a single consumption good and are therefore silent about the optimal degree of rate differentiation between different consumption categories. The work of Conesa et al. (2020) features two types of consumption goods and households with non-homothetic preferences. They study the welfare implications of replacing the current US tax system by a consumption tax and show that in the absence of all other tax instruments, the optimal tax on basic goods is markedly lower than on luxury goods. While their focus is on replacing one system with another, I emphasize the interaction and joint-design of different tax instruments. Finally, I consider a substantially larger number of consumption categories. Parodi (2023) analyzes the optimal joint-design of consumption and income taxes using a life-cycle model calibrated to Italy. She shows that durables should be subsidized and consumption goods taxed with a uniform tax. My focus, instead, is on the optimal tax on different consumption goods rather than the relationship to taxes on durables. Finally, I do not *ex ante* categorize goods into luxuries and necessities but my model endogenously categorizes goods and services into necessities and luxuries.

My work is also related to the Public Finance literature studying the design of consumption taxes. Early work on the topic was primarily focused on second-best questions of how to design consumption taxes when other instruments are not available. For instance, Ramsey (1927) argued that optimal commodity taxation is directly related to a good's price elasticity. Starting with the seminal results by Atkinson and Stiglitz (1976) and Diamond and Mirrlees (1971), extensive work has been done to better understand the situations in which non-uniform consumption taxes are optimal in a first-best setting. The two papers combined basically argue that no differentiated commodity taxation is required when an optimal income tax is available since this avoids distortions to production and consumption choices. These results continue to be of great importance but have been qualified by various more recent work. For instance, Saez (2002) or Golosov et al. (2013) showed that heterogeneous consumption tax rates can be optimal when households have heterogeneous preferences. Similarly, non-uniform rates have been found to be optimal when the consumption of certain goods and services creates externalities (Bastani and Köhne, 2022). My paper, instead, analyzes the question at stake from a quantitative perspective capturing many of the empirically relevant trade-offs related to the optimal tax design. My model generates income, wealth, and consumption inequality which is in line with what we observe in the data and my optimal tax design is therefore affected by all these different dimensions of heterogeneity. Furthermore, I can make statements about the optimal tax on a variety of goods and services present in US household's everyday consumption baskets. Finally, the fact that my model includes an explicit intertemporal consumption-savings decision allows me to draw conclusions about the dynamic effects of consumption taxes and their effects on savings and capital accumulation.

Outline. The structure of the paper is as follows. I first describe the quantitative model used to analyze the optimal design of consumption taxes. In Section 3, I bring the model to the data and describe the estimation and calibration procedure. Section 4 solves for the optimal joint design of consumption taxes and labor income taxes. Finally, Section 5 discusses various factors resulting in the optimal tax rates found in the previous section. The last section concludes.

2 Quantitative Model

In this section, I develop a quantitative general equilibrium model to analyze the optimal design of consumption taxes. Its basic structure follows incomplete markets model in the spirit of Aiyagari (1994), Huggett (1993), or Bewley (1983) with several crucial deviations: Households consume a variety of consumption goods and services, over which they have non-homothetic preferences. They endogenously decide on the number of hours worked, and a government sector raises revenue through various tax instruments and uses the proceedings to finance government consumption, to make targeted transfers to households, and to repay debt.

Environment. Time is discrete and infinite with periods $t = 0, 1, ...$ There are heterogeneous households, a government, and firms producing a variety of consumption goods. Factor and goods markets are competitive and both firms and households are atomistically small and take prices as given. The asset market is incomplete and households can only save in a one-period risk-less asset.

Households. The economy is populated by a continuum of measure one of infitely-lived households, which differ in their asset holdings a, their idiosyncratic productivity z and their skill type s. Asset holdings are a choice of the households. Idiosyncratic productivity is assumed to be stochastic given an exogenous process with some persistence. At the beginning of time, there is a probability p_s that a household is high-skilled and with probability $1 - p_s$ the household is low-skilled. This skill-type is permanent, and at every point in time, a fraction p_s is of the high-skilled type.

Recursive Household Problem. In period t, an agent of skill type s with productivity shock and asset level (z_t, a_t) solves:

$$V_t(z_t, a_t, s) = \max_{E_t, n_t, a_{t+1}} \left\{ U(E_t) - v(n_t) + \beta \sum_{z \in \mathcal{Z}} \Pi(z_{t+1} \mid z_t) V_{t+1}(z_{t+1}, a_{t+1}, s) \right\}$$

subject to

$$\underbrace{\underbrace{\sum_{j=1}^{J} c_{j,t} p_{j,t}(1+\tau_{j,t})}_{\equiv E_{\mathcal{C}}} + \underbrace{c_{o,t} p_{o,t}}_{\equiv E_o}}_{\equiv E} + a_{t+1} \leq w_t z_t s n_t - \mathcal{T}_t(w_t z_t s n_t) + (1 + r_t(1-\tau_{k,t})) a_t$$

$$c_{j,t}, c_{o,t}, a_{t+1} \geq 0 \; \forall j \text{ and } n_t \in [0,1]$$

In every period, households decide how many hours to work n_t, how much to save using the riskless asset a_{t+1} and how much to allocate to expenditures E_t. Utility from expenditures is denoted $U(E_t)$ and the disutility from working $v(n_t)$. $\beta \in (0,1)$ is the subjective discount factor and $\Pi(\cdot|\cdot)$ the transition matrix for household productivity z. Households earn labor income $wzsn$ depending on the wage w, their household productivity z, their skill-type s[1], and the hours worked n. Labor income is taxed with the non-linear tax function $\mathcal{T}(\cdot)$ and capital income is taxed at a flat rate τ_k.
There are two types of expenditures: pure consumption expenditures $E_{\mathcal{C}}$ and expenditures on the outside good E_o, which sum up to total expenditures E. $E_{\mathcal{C}}$ summarizes the expenditures on pure consumption goods such as food at home, food in restaurants, personal services or recreational expenses. Consumption goods[2] are indexed by $j \in J$, priced at price p_j and potentially taxed at rate τ_j.
Empirically, households do not spend all their income net of savings on consumption, but also pay for rent and mortgages, or for insurance coverage. The model explicitly accounts for these "outside expenditures" E_o, which are not subject to the sales tax. As will be shown when calibrating the model, the split between taxed consumption goods and outside expenditures is non-homothetic in itself when looking at US consumption data. This is crucial for the optimal tax design since low-income households spend a larger share of their total expenditures on (potentially) taxed consumption goods. For the analysis that follows, the outside expenditures are explicitly not applied a consumption tax since this is the current empirical pattern observed in the US. The analysis in this paper taxes the tax system as given and optimizes over the existing set of tax instruments. Further analysis could take into account what the implications are of applying the consumption tax to a wider set of expendtitures, such as rent or mortgage payments.
Households engage in a multi-stage budgeting behavior. First, they face a simultaneous labor choice and intertemporal savings-expenditure decision given $U(E)$ and $v(n) = B_l \frac{n^{1+\chi}}{1+\chi}$. This de-

[1] Note that the skill-type is permanent and I therefore drop any time subscripts, both in the text and in the exposition of the model.

[2] If not specified otherwise, I use the term "consumption goods" for both consumption goods and services.

cision implies a level of total expenditures E, which is split into consumption expenditures and outside expenditures given the functional form of $U(E)$:

$$U(E) = \frac{[\mathcal{C}(E_{\mathcal{C}})^{\eta}(C_o(E_o) - \underline{C}_o)^{1-\eta}]^{1-\theta}}{1-\theta}$$

where $E = E_{\mathcal{C}} + E_o$. For $\underline{C}_o \neq 0$, the split between $E_{\mathcal{C}}$ and E_o is non-homothetic. In the data, high income households allocate a higher share of their expenditures to the outside good than low-income households and $\underline{C}_o$ will therefore be negative. The final decision concerns the split of $E_{\mathcal{C}}$ into the J consumption goods and services. $\mathcal{C}$ is the non-homothetic Constant Elasticity of Substitution (CES) aggregator that aggregates the individual consumption goods c_j. The CES aggregator is implicitly defined following Comin et al. (2021):

$$\sum_{j=1}^{J} \Upsilon_j^{\frac{1}{\sigma}} \left(\frac{c_j}{\mathcal{C}^{\psi_j}} \right)^{\frac{\sigma-1}{\sigma}} = 1 \tag{1}$$

Υ_j is the share parameter for good j, σ the (constant) elasticity of substitution between the different consumption goods, and ψ_j introduces the non-homothetic pattern of consumption across the various consumption goods. ψ_j also governs the expenditure elasticity of each good. This functional form is particularly suited for the analysis of the question asked in this paper. It easily accommodates an arbitrary number of goods, whose non-homotheticity is governed by a single parameter ψ_j. I will also show in Section 3 that the preferences can be brought to the data, thereby providing credible estimates of the parameters ψ_j, Υ_j and σ. Lastly, as compared to other classes of non-homothetic preferences, such as generalized Stone-Geary preferences, the slopes of the Engel curves implied by non-homothetic CES preferences do not flatten as income grows and therefore allow for a realistic consumption behavior – even with very high levels of income. The allocation of consumption expenditures can be characterized by a few equations. Given an overall level of $E_{\mathcal{C}}$, c_j is determined by:[3]

$$c_j = \Upsilon_j \left(\frac{p_j(1+\tau_j)}{E_{\mathcal{C}}} \right)^{-\sigma} \mathcal{C}^{(1-\sigma)\psi_j} \tag{2}$$

Hence, given $\sigma > 0$, the demand for good j is decreasing in the prices faced by consumers ($p_j(1+\tau_j)$) and increasing in the share parameter Υ_j. As total expenditures increase, so does consumption of good c_j but the degree to which it does so depends on the change through $\mathcal{C}$. $E_{\mathcal{C}}$ and $\mathcal{C}$ are implicitly connected though the following expression:

$$E_{\mathcal{C}} = \left(\sum_{j=1}^{J} (\mathcal{C}^{\psi_j} p_j (1+\tau_j)^{1-\sigma}) \right)^{\frac{1}{1-\sigma}} \tag{3}$$

Taken together, Equations 1-3 fully characterize the consumption allocation of households across the different consumption goods. Hence, for $\sigma > 0$, the aggregator $\mathcal{C}$ is increasing in total con-

[3] The allocation of expenditures within the J consumption categories follows the work of Comin et al. (2021) and the derivations of these expressions can be found in their appendix.

sumption expenditures. Compared to the homothetic case ($\psi_j = 1\,\forall\, j$), the relationship between expenditures and consumption shares is, however, non-linear. To make visible how exactly the non-homotheticity parameters ψ_j determine relative consumption behavior, one can derive a decomposition of the relative expenditures allocated to each good as a function of price effects and income (expenditure) effects. Using Equation 1, one can show that:

$$\log\left(\frac{\omega_{j,t}}{\omega_{k,t}}\right) = (1-\sigma)\log\left(\frac{p_{j,t}(1+\tau_{j,t})}{p_{k,t}(1+\tau_{k,t})}\right) \tag{4}$$

$$+ (1-\sigma)\left(\psi_j - \psi_k\right)\underbrace{\left[\log\left(\frac{E_{\mathcal{C}t}}{p_{b,t}}\right) + \frac{1}{1-\sigma}\log\omega_{b,t}\right]}_{\equiv \log \mathcal{C}_t} + \log\left(\frac{\Upsilon_j}{\Upsilon_k}\right) \tag{5}$$

$\left(\frac{\omega_{jt}}{\omega_{kt}}\right)$ is the relative expenditure share allocated to good j at time t as compared to good k. $\omega_{b,t}$ denotes the share allocated to a base good. The first term on the right-hand side represents the contribution of prices and taxes to the determination of expenditure shares and acts equivalently to the case of homothetic preferences. The second term, instead, highlights the effect of the non-homotheticity: Changes in the total amount of consumption expenditures $E_{\mathcal{C}}$ affect the relative expenditure shares depending on the relative size of the non-homotheticity parameters ψ_j and ψ_k. If $\psi_j > \psi_k$, then increases in total consumption expenditures $E_{\mathcal{C},t}$ increase the relative share allocated to good j as compared to good k.

Government. The government in this economy collects a variety of taxes: Consumption taxes $\{\tau_{j,t}\}_{t=0}^{\infty}$ for $j = 1, ..., J$, capital income taxes $\{\tau_{k,t}\}_{t=0}^{\infty}$ and non-linear labor income taxes $\{\mathcal{T}_t(y_l)\}_{t=0}^{\infty}$, where $y_l = wzsn$. I assume the labor income tax to take the log-linear form as commonly used in the quantitative public-finance literature (see Bénabou (2002) or Heathcote et al. (2017)):

$$\mathcal{T}(y) = y - \lambda(y)^{1-\tau} + T$$

$1-\lambda$ is a measure of the average labor income tax[4] and τ the progressivity of the tax code. Absent the lump-sum transfer and whenever $\tau > 0$, the labor income taxes are progressive. Depending on the level and progressivity of the tax code, some households may pay negative net taxes (i.e. receive transfers). To accurately match the degree of redistribution in the US, the standard log-linear tax function needs to be supplemented by a lump-sum tax T (Boar and Midrigan, 2022). The government also issues and repays debt D. I assume that the government does so on international capital markets and that the interest paid on government debt corresponds to the domestic equilibrium interest rate. Finally, the government finances an exogenous stream of government expenditures G. The government budget constraint is required to hold every period:

[4]Strictly speaking, λ shifts the tax function up or down and determines the average level of taxation in the economy. It should, however, not be understood as the average tax rate. Furthermore, looking at the tax function without the lump-sum transfer, $y_0 = \lambda^{\frac{1}{\tau}}$ implicitly defines the break-even income level (y_0) where households with income above y_0 pay taxes and households below y_0 recive net-transfers (Heathcote et al., 2017).

$$G + (1+r)D = D + \underbrace{\int T(y)d\mu(\cdot)}_{\text{Labor income tax}} + \underbrace{\sum_{j=1}^{J} \tau_j \int c_j(a,z,s)d\mu(\cdot))}_{\text{Consumption tax}} + \underbrace{r\tau_k \int a(a,z,s)d\mu(\cdot))}_{\text{Capital income tax}}$$

The left-hand side summarizes government expenditures, and the right-hand side shows government revenue in steady state. $\mu(\cdot) = \mu(a,z,s)$ denotes the measure of households with state (a,z,s) and $c_j(\cdot), n(\cdot), c_o(\cdot)$ and $a(\cdot)$ are the optimal policies resulting from household optimization.

Firms. There is one intermediate goods firm, which uses capital and labor as inputs and produces the intermediate good Y_0. Y_0 constitutes the numeraire in the economy. This firm solves the standard maximization problem:

$$\max_{K,L} AK^{\alpha}L^{1-\alpha} - wL - (R+\delta)K$$

where δ is the depreciation rate, w the wage, and $R = 1 + r$ the gross rental rate of capital. The basic good can also be considered the outside good and households can consume the outside good at price $p_o = 1$. There are also J specialized firm that use the basic good to transform it into the different consumption goods using the linear transformation technology:

$$Y_j = \gamma_j Y_0$$

$\forall\, j = \{1, 2, ..., J\}$. Hence, the price for each consumption good is $p_j = \frac{1}{\gamma_j}$. The equilibrium in this economy is defined in Appendix A.1.

3 Quantifying the Model

There is a large number of parameters that need to be estimated or calibrated. I, therefore, split the estimation procedure into three steps in order to bring the model to the data. A limited number of parameters is externally calibrated using established estimates from the literature. I then estimate the parameters of the non-homothetic CES preferences using US household data. Finally, I estimate the remaining parameters to match various moments observed in the data. This separate estimation of the non-homothetic CES and the remaining parameters is possible thanks to the weak separability between consumption and outside expenditures as well as the full separability between total expenditures and labor supply (Parodi, 2023). One period in the model corresponds to one quarter.

Estimation non-homothetic CES

I estimate the preferences on US household data using the Generalized Method of Moments (GMM). Following, I will briefly describe the dataset as well as the estimation strategy, which follows very closely Comin et al. (2020) and Comin et al. (2021).

Data. To estimate the non-homothetic preferences, I rely on consumption data from the US Consumption Expenditure Survey (CEX) covering the years 2000 to 2019. The CEX is a stratified survey conducted by the U.S. Bureau of Labor Statistics (BLS). Households are interviewed in 4 consecutive quarters[5] and the sample includes approximately 6000 households per quarter (Borusyak and Jaravel, 2022). The CEX consists of two distinct surveys: the diary survey, which collects minor and more frequent purchases and the interview survey, which contains major and recurring items. For the estimation, I rely on the latter to capture a wide range of expenditure items. The interview survey contains detailed information on hundreds of expenditure items purchased over a period of three months preceding the interview. I manually aggregate these items into 12 overarching expenditure categories: one category capturing all outside expenditures and 11 consumption expenditure categories.[6] The matching into distinct categories is conducted such that there is maximum homogeneity within groups and maximum heterogeneity across groups and to ensure consistency across CEX waves. During the last interview, respondents are also asked questions about their income situation and various demographic characteristics. Whenever I use income, I rely on pre-tax income, represented in 2019 USD. In selecting the sample, I closely follow Aguiar and Bils (2015) and Comin et al. (2021). First, I restrict my sample to include only urban households between the age of 25 and 64. Second, I drop households with zero expenditures on the *Food at home* category, to avoid issues related to misreporting of expenditures. Finally, I follow Borusyak and Jaravel (2022) and drop households with income below $5,000 due to concerns about misreporting and temporary unemployment. Finally, since my estimation strategy makes use of the panel structure of the data, I only use households which are observed over 4 consecutive quarters.

I combine the CEX data with regional price data from the BLS' urban CPI. I use prices for various expenditure categories on the census level, which divides the US into four regions (Northeast, Midwest, South, and West). To construct a household-specific price measure for each of the 11 categories considered, I compute an expenditure-weighted average of the log-price of each of the sub-categories entering the 11 categories considered. For instance, the category *Recreation and Sports* consists of sports expenditures, holiday expenditures and book expenditures and I therefore use the corresponding regional price indices for each of these sub-categories and the household expenditures weights to compute a household-specific price index of *Recreation and Sports*.

Estimation Strategy. Making use of the log-linear shape of Equation 5 and using the dataset described above, I can estimate the parameters of the non-homothetic CES preferences. Starting from Equation 5 and by rearranging, I obtain an empirical representation of the relative demand for good j as compared to the base good b:

[5]Until 2015, 5 interviews were conducted where data from 4 interviews was available for public use.

[6]The 11 categories are: *Food at home, Food away from home, Alcohol and Tobacco, Utilities, (Public) Transportation, Vehicle Expenses, Personal Services, Domestic Services, Recreation and Sports, Entertainment, Housing Material and Communication.*

$$\log\left(\frac{\omega^n_{j,t}}{\omega^n_{b,t}}\right) = (1-\sigma)\log\left(\frac{p^n_{j,t}(1+\tau^n_{j,t})}{p^n_{b,t}(1+\tau^n_{b,t})}\right) + (1-\sigma)\left(\psi_j - 1\right)\log\left(\frac{E^n_{C,t}}{p^n_{b,t}(1+\tau^n_{b,t})}\right) + \left(\psi_j - 1\right)\log\omega^n_{b,t} + \xi^n_j + \epsilon^n_{j,t} \tag{6}$$

As earlier, the left-hand side is the logarithm of the ratio of the expenditure share that household n allocates to good j at time t relative to the share allocated to the base good b. $p^n_{j,t}(1+\tau^n_{j,t})$ is the price faced by consumer n for good j at time t. Prices and taxes can potentially vary at the household level depending on the Census region, in which the household lives and since I use household-level expenditure shares to create household-specific price indices for the consumption categories. As earlier, a base category needs to be chosen, for which the ψ_j is normalized to 1. I use *Domestic Services* for this. The estimation results are independent of the exact choice of the base category, but I choose *Domestic Services* as it will turn out to be a category with an intermediate level of non-homotheticity.

ξ^n_j corresponds to the relative size of the preference parameters: $log\left(\frac{Y^n_j}{Y^n_b}\right)$. In line with Comin et al. (2021), I assume that these preference parameters are fully explained by a selection of household characteristics[7] X^n and region-category fixed-effects $\mu_{j,r}$: $\xi^n_j = \beta'_j X^n + \mu_{j,r}$. Finally, the error term is assumed to be the composite of category-time fixed-effects $\mu_{j,t}$ and a mean-zero error term $\tilde{\epsilon}^n_{jt}$: $\epsilon^n_{j,t} = \mu_{j,t} + \tilde{\epsilon}^n_{jt}$. There is one such equation for each of the $J-1$ categories, which can be used to simultaneously estimate the non-homotheticity parameters and the elasticity of substitution. Controlling for household characteristics (X^n), this estimation strategy exploits within-region co-variation between expenditure shares and total household expenditures to identify the parameters of interest (Comin et al., 2021).

Measurement error and endogeneity issues can bias the results of estimating equation 6. Aguiar and Bils (2015) highlight that small measurement errors at the expenditure item level can result in large measurement errors if aggregated up to construct $E^n_{C,t}$. To address this issue, I instrument total consumption expenditures using total household income in the preceding year. In the CEX, this is a single question and does, therefore, not suffer from the aggregation of numerous small measurement errors. Second, regional shocks can induce endogeneity issues with respect to the price indices used. To mitigate the bias resulting from this, I instrument relative prices using a Hausman price instrument (Comin et al., 2021). The idea of the instrument is to produce a price series for household n and good j, which captures common trends in US prices and is free of measurement error at the household level. To do so, I compute a price for household n, good j at time t by taking the average price in all regions except the one in which household n resides, and use the average expenditure shares in the residence region of the household to construct the household-specific price instrument. Using all prices except those from the own region alleviates the endogeneity arising from regional shocks, and using aggregate expenditure shares addresses the measurement issues under the assumption that the measurement errors average out within a

[7] I include dummies for different age bins and the number of earners in the household.

region.
With this at hand, the system of $J-1$ equations induced by expression 6 can be estimated using the generalized method of moments (GMM). I use a 2-step GMM procedure and cluster standard errors at the household level. Furthermore, I impose σ to be identical across all equations.

Results. Table 1 presents the estimation results. The second column displays the estimated parameters, $\{\psi_j\}_{j=1}^{J}$ and σ. The last two columns summarize the implied expenditure elasticities for the first and 5^{th} income quintile. With the type of preferences used, expenditure elasticities depend on the total level of consumption expenditures E_C and are therefore income-dependent.[8]

TABLE 1: GMM Estimation Results

Variable	**Estimates**	$\eta_{j,1}$	$\eta_{j,5}$
σ	0.16 (0.03)		
ψ: **Food at home**	0.16 (0.01)	0.43	0.36
ψ: **Utilities**	0.36 (0.01)	0.77	0.61
ψ: **Housing Material & Communication**	0.48 (0.01)	0.98	0.76
ψ: **Public Transportation**	0.61 (0.02)	1.20	0.92
ψ: **Alcohol & Tobacco**	0.71 (0.02)	1.37	1.05
ψ: **Entertainment**	0.96 (0.02)	1.81	1.37
ψ: **Domestic Services**	1.00 (–)	1.87	1.42
ψ: **Personal Services**	1.31 (0.02)	2.41	1.81
ψ: **Vehicle Expenses**	1.32 (0.03)	2.42	1.82
ψ: **Food away from home**	1.49 (0.03)	2.71	2.03
ψ: **Recreation & Sports**	2.01 (0.04)	3.61	2.69

Comments: Estimated using GMM and CEX (2000-2019),
79'800 obs., 19'950 households, clustered SE in parentheses (clustered at household level)
$\eta_{j,p}$: Expenditure elasticity for p-th income quintile implied by the estimates
Domestic Services constitutes the base category and its ψ is normalized to 1

My estimate for the elasticity of substitution (σ) is 0.16, which implies that the 11 expenditure categories are complementary to each other. This value is slightly lower than what Comin et al. (2020) find for 8 expenditure categories but has the same qualitative conclusion regarding the substitutability of expenditure categories.[9] While the ψ parameters do not have a direct meaning, a higher value generally speaks for a more luxurious good with a higher expenditure elasticity. The non-homotheticity parameters $\{\psi_j\}_{j\in J}$ vary from relatively low values of 0.16 for *Food at home* to 2.01 (*Recreation and Sports*). The implied expenditure elasticities imply that 3 out of the 11 goods are considered necessity goods for households in the first income quintile whereas richer households consider 4 categories as necessities. *Public Transportation* is a necessity for high-income households but a luxury good for low-income households.

[8] The corresponding expression is $\eta_{j,q} = \sigma + (1-\sigma)\frac{\psi_j}{\bar{\psi}}$ where σ is the elasticity of substitution and $\bar{\psi}$ the expenditure weighted average of all ψ_j. See Comin et al. (2021) for a proof.

[9] Contrary to this study, Comin et al. (2020) also do not focus on pure consumption items and estimate the preferences on value-added consumption which can further explain the slight difference in the parameter estimate.

Calibration

Having estimated the parameters governing household consumption behavior using GMM and US household data, I now calibrate the remaining parameters to match various US data moments and set a limited number of parameters according to established values from the literature.

Share Parameters. In a first step, I follow Comin et al. (2020) and set the share parameter $\{\Upsilon_j\}_{j\in J}$ to match the average expenditure share observed in the data. In the previous GMM estimation, I assumed that the share parameters are fully explained by a set of household characteristics and fixed effects. Hence, if a majority of the sample comes from a specific age group or has a particular number of earners, then the estimated $\{\Upsilon_j\}_{j\in J}$ should deviate from 1 in order for the model to be able to match the average share observed in the data. Table 7 in the Appendix summarizes the results. The share parameters range from 0.15 for *Public Transportation* to 2.22 for *Food at home*.

Earnings Process. Accurately representing the labor earnings process is crucial for the analysis of tax systems, as it directly determines households' motivation to self-insure against idiosyncratic risk. To capture the rich dynamics of labor earnings, I assume that households' labor productivity follows a Gaussian Mixture Autoregressive (GMAR) process in logs:[10]

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

With probability $1-p$, households draw an innovation from a normal distribution with low variance (σ_2^2). Infrequently (with probability p), a high-variance innovation is drawn. I use the values estimated by Nord et al. (2025), who target the (i) the cross-sectional variance of log annual earnings, (ii) the standard deviation, (iii) the skewness and (iv) kurtosis of log annual earnings changes, and the (v) ratio of the 90^{th} to the 10^{th} percentile of log changes of annual earnings changes as derived from the PSID by Nardi et al. (2019).[11] The estimation procedure implies $\rho = 0.963$, $\sigma_1 = 0.5$, $\sigma_2 = 0.91$ and $p = 0.156$. I use the algorithm provided by Farmer and Toda (2017) to discretize the GMAR process.

Households and Firms. In addition to the income process and the parameters defining the non-homothetic CES, there are various other parameters related to the household that I estimate by targeting specific data moments or where I use established values from the literature. Table 2 summarizes these parameters as well as the ones relating to the firm problem. The top panel collects

[10] See, for instance, Ferriere et al. (2023) or Nord et al. (2025) for applications of this income process to heterogenous agents models.

[11] The dataset includes PSID waves for the years ranging from 1962 to 1992. The sample includes only households with household heads aged 25 to 60, and household-level earnings are adjusted for year fixed effects and family size (Nord et al., 2025; Nardi et al., 2019).

TABLE 2: Model Calibration

Target	Model	Data	Closest Parameter	Source
Internally calibrated				
$\frac{K}{Y}$ Ratio	3.2	3.0	$\beta = 0.958$	Penn World Tables
Consumption Expenditures Q1 $E_{\chi,Q1}$	4.75	4.46	$\phi = 2.72$, $\underline{C}_o = -5.36$	CEX 2001-2019
Consumption Expenditures Q5 $E_{\chi,Q5}$	9.20	9.33	$\phi = 2.72$, $\underline{C}_o = -5.36$	CEX 2001-2019
Consumption Ratio Q1 $\frac{E_{\chi,Q1}}{E_{Q1}}$	0.39	0.39	$\eta = 0.26$, $\underline{C}_o = -5.36$	CEX 2001-2019
Consumption Ratio Q2 $\frac{E_{\chi,Q2}}{E_{Q2}}$	0.34	0.27	$\eta = 0.26$, $\underline{C}_o = -5.36$	CEX 2001-2019
Hours worked n	$\frac{1}{3}$	0.31	$B_l = 9.46$	Ferriere et al. (2023)
Mean Expenditure Share on good j $\frac{c_j p_j (1+\tau_j)}{E_c}$			Υ_j, see Table 7	CEX 2001-2019
Difference Post-tax income and Pre-tax income Q1	-62%	-62%	$\tau = 0.12$, $T = 11.4\%$ of med. income	CBO 2019
Difference Post-tax income and Pre-tax income Q5	29%	24%	$\tau = 0.12$, $T = 11.4\%$ of med. income	CBO 2019
Hours worked n	$\frac{1}{3}$	0.31	$B_l = 9.46$	Ferriere et al. (2023)
Externally calibrated				
Risk Aversion θ	2			Comin et al. (2021)
Permanent Productivity Low-Skilled s_1	1			normalization
Permanent Productivity High-Skilled s_2	1.8			CPS 2018
Share High-Skilled ω_s	0.37			CPS 2018
Inverse Frisch Elasticity χ	2.5			Ferriere et al. (2023)
Capital Share α	0.33			standard
Transformation Efficiency $\gamma_j = 1 \ \forall j$	1			see text
Depreciation Rate δ	0.015			Ferriere et al. (2023)

the parameters, which are determined internally. Changes in parameters affect a variety of model moments simultaneously and I therefore jointly solve for the parameter values that match best the data. Nonetheless, there are parameters, which are more closely related to certain data moments. The parameter closest to a specific moment is listed in the *Closest Parameter* column of the table.

I target a capital-income ratio of roughly 3.0 by adjusting the subjective discount factor β. Since I work with non-homothetic preferences, absolute expenditures matter for the composition of the consumption basket. I therefore use the scaling parameter ϕ and the subsistence level of the outside good to target the amount of expenditures allocated to consumption expenditures of the first and the 5^{th} income quintile in 1000 USD. To accurately capture the fact that low and high-income households allocate different shares of their total expenditures to consumption goods, I calibrate η and the subsistence level of the outside good $\underline{C}_o$ to jointly match the share of expenditures allocated to consumption expenditures of the first and second income quintile. This ensures that both the level and the slope across the income distribution in the model match their empirical counterparts. In the data, the share of outside expenditures is monotonically increasing in income and targeting only the first and second income quintiles, therefore ensures that the pattern is consistent for the entire income distribution. Hours worked are matched by targeting $\frac{1}{3}$ using the parameter determining disutility from work (B_l). 8 hours of work correspond to exactly $\frac{1}{3}$ of 24 hours. Finally and as described above, the mean expenditure share on each consumption good is matched using the share parameters $\{\Upsilon_j\}_{j \in J}$.

The lower part of the table summarizes the parameters, which are set according to values from the literature. As in Comin et al. (2021), I set the risk aversion parameter θ to 2. I normalize the productivity of low-skilled workers to 1 and set s_2 to 1.8 to match the college premium in the US in 2018 as estimated by the ratio of the median weekly earnings of bachelor degree holders relative to high school graduates without any further degree.[12] Similarly, the share of high-skilled workers, 0.37, is drawn from the Current Population Survey for the year 2018.[13] I take the value for the Inverse Frisch Elasticity from Ferriere et al. (2023). Finally, on the firm side, I set the capital share to 0.33 as commonly assumed in the literature. I do not assume that there are any differences in the transformation technology used by the different firms and set all $\{\gamma\}_{j\in J}$ to 1. Finally, the depreciation rate δ is set to 1.534% to obtain an annual depreciation rate of 6% as in Ferriere et al. (2023).

Government. The last block of parameters govern the government sector. As listed in Table 2, I use the lump-sump transfer T and the progressivity of the labor income tax τ to target the percentage different between pre- and posttax income of the first and 5th income quintile in the CBO data.[14] With the exception of these two parameters, all parameters are calibrated externally. The capital income tax is set to 36% as suggested by Trabandt and Uhlig (2011). Government spending as a fraction of GDP and the debt-to-GDP ratio are set to 21.8% and 106.4% respectively following Ferriere et al. (2023). US states and counties have discretion over their sales taxes and, consequently, they vary across states and counties. However, most states' sales taxes are within the range of 0% to 9%. Trabandt and Uhlig (2011) estimate the average US sales tax to be 5%, which I assume to be the *status quo* value of consumption taxes in the US for all expenditure categories. Outside expenditures capture a variety of expenditures, such as monetary contributions, rental payments, mortgage payments, or insurance premiums. These expenditure categories are not subject to the sales tax in the US and I therefore do not apply a tax to the outside good. Finally, I internally calibrate the lump-sum tax by targeting the average tax rate paid by the first income quintile, which is -7% (Ferriere et al., 2023).

TABLE 3: Calibration of Government Parameters

Parameter	Description	Value	Source
τ_k	Capital Income Tax	0.36	Trabandt and Uhlig (2011)
g	Government Spending G/Y	0.218	Ferriere et al. (2023)
d	Public Debt D/Y	1.064	Ferriere et al. (2023)
τ_n	Consumption Tax Necessities	0%	reduced rate in most US states
τ_l	Consumption Tax Luxuries	5%	Trabandt and Uhlig (2011)
τ_o	Tax Outside Good	2.5%	see text
T	Lump Sum Tax	11.43% of median income	estimated
τ	Progressivity Income Tax	0.12	estimated

3.1 Model Fit

I next analyze how the model compares to matched and unmatched data moments.

[12] I use data from the Current Population Survey of the year 2018, as presented in the FRED Blog (FRED, 2018).

[13] See US Census Bureau (2023).

[14] I use data from the CBO's overview on the distribution of income for the year 2019 (Congressional Budget Office, 2022).

Inequality. Figure 1 compares various measures of inequality in the model with their equivalent moments in the data. I targeted various data moments related to labor earnings when calibrating the stochastic productivity process but the data moments in Figure 1 are not directly targeted. From the left to the right, it displays inequality of labor income, wealth, and consumption expenditures.

FIGURE 1: Fit of Inequality Measures

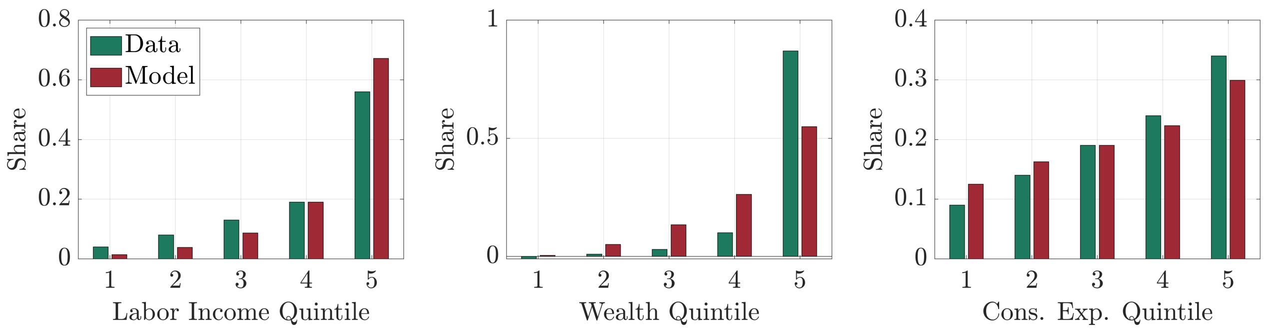

Note: This figure plots various measures related to income, wealth and consumption inequality in the data and in the model economy. Data on *Labor Income* and *Wealth* originates from the SCF 2013 (Federal Reserve Board, 2013). Data on *Consumption Expenditures* comes from the CEX, 2001-2019 (Bureau of Labor Statistics, 2020).

Despite not being targeted, the model does a good job at matching the different dimensions of inequality measured in the data. While high labor income households earn a slightly too large share of total income, the overall pattern is clearly in line with what is observed in the data. Conversely, the model is not able to match the vast inequality in wealth at the top of the income distribution. While in the data the 5^{th} wealth quintile holds more than 80% of total wealth, the share is only slightly above 50% in the model. Conversely, the share held by the lower 80% of the wealth distribution is slightly too high. Nonetheless, given that my model does not include any added extensions to be able to match the vast inequality in wealth observed in the data, the benchmark version provides a reasonably accurate representation of the wealth inequality observed in the data.

Crucially, the model does a good job at matching the degree of consumption expenditure inequality in the data. While the the slope of the share of consumption expenditures across the distribution is slightly to flat in the model as compared to the data, the overall picture in the model mirrors what is observed in the Consumption Expenditure Survey.[15] Overall, however, the model matches the different dimensions of inequality in the data reasonably well.

Figure 2 displays the fit of the tax and transfer system implemented in the economy alongside the consumption taxes. It shows the percentage difference between pre and posttax income for both the data and the model. Overall, the model does a good job at matching the degree of redistribution taking place in the US, with all differences very close to 0.

Finally, the model generates a dispersion of hours slightly lower than what is observed in the data

[15] Table 8 in the appendix provides additional summary statistics related to earnings and consumption inequality. The conclusions remain the same even for more detailed statistics.

FIGURE 2: Fit of the Tax and Transfer Rate

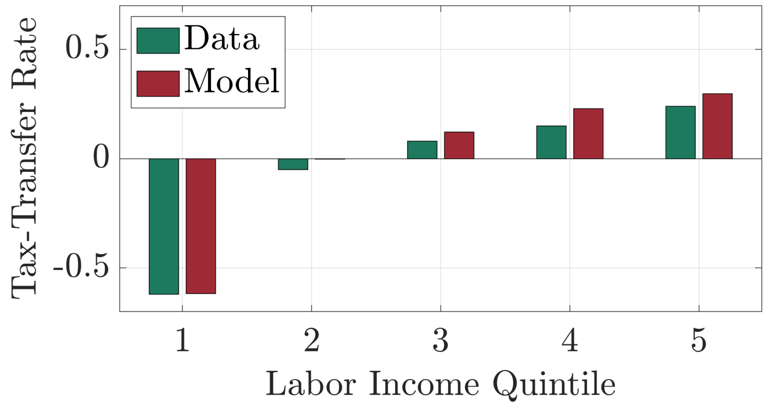

Note: This figure shows the percentage difference between pre and post-tax income in the data and in the model economy. The data comes from the Congressional Budget Office (2022) and covers the year 2019.

but consistent with similar models.[16] Whereas, my model generates a variance of log-hours of roughly 0.04, the empirical value is around 0.09. (Ferriere et al., 2023).

Heterogeneous Consumption Baskets. The various plots in Figure 3 display that the model replicates well the consumption behavior of US households. The absolute amount of income allocated to consumption goods (E_C) matches what is observed in US household data (plot in the top left corner): The higher total household income, the higher the absolute amount allocated to consumption expenditures. While low-income households spend around 4'500 USD per quarter on consumption goods and services, this amount rises to more than 9'000 USD for households in the 5^{th} income quintile. This indicates that the model captures well the combined decision along two margins: how much to allocate to savings versus expenditures and how much to allocate to consumption expenditures versus the outside good. Note that I only target the first and 5^{th} income quintiles' values, and all other values are unmatched. There is a slight deviation of the model from the data for the 2^{nd} income quintile but for all quintiles, the margin of error is below 100USD.
A key argument of this paper is that households differ in the allocation of their consumption expenditures across different consumption goods and services and that this matters for the optimal tax design. The remaining 11 plots in Figure 3 summarize the expenditure shares allocated to each of the 11 consumption expenditure categories as a function of total household income. Recall that the parameters determining the shape of expenditure patterns originate from the GMM estimation of the $J - 1$ demand equations and are therefore not directly targeted moments. The level of expenditure shares was estimated in a second step by targeting the average (economy-wide) expenditure shares allocated to good j and is therefore directly targeted. Overall, the plots show that the non-homothetic CES preferences allow for a realistic representation of the consumption behavior across the income distribution in the US.

[16] The model of Ferriere et al. (2023) generates a dispersion of 0.06

FIGURE 3: Fit of Consumption Baskets

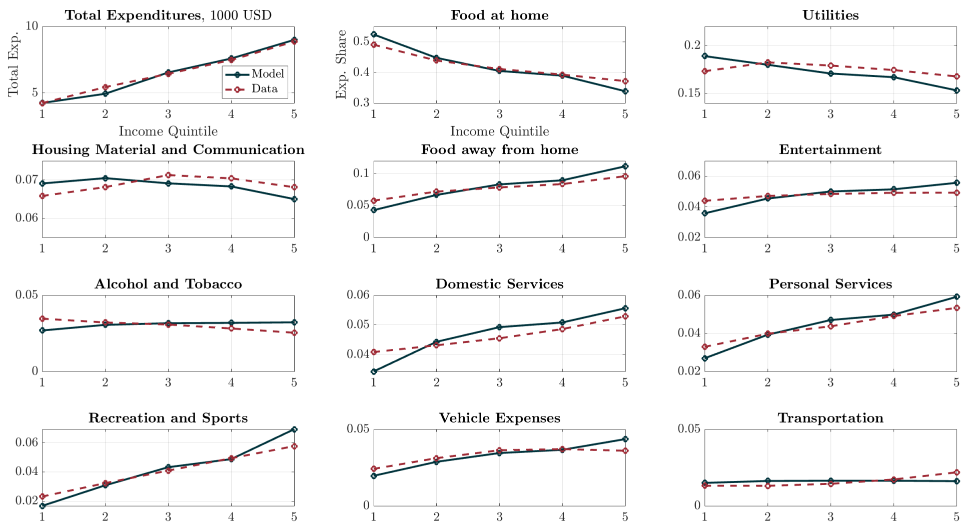

Note: The figure shows consumption across the income distribution in the data and in the model economy. Each dot stands for an income quintile and the lines stand for the expenditure share allocated to a given category. The exception is the first plot in the top right corner, which shows the level of total expenditures across the income distribution. Data comes from the CEX, covering the year 2001 to 2019 (Bureau of Labor Statistics, 2020).

Low-income households allocate around 50% of their consumption budget to groceries (*Food at home*). This share is still around 35% for high-income households. *Utilities* also constitute a large share of households' consumption expenditures and, similar to groceries, the share is clearly decreasing in income. *Housing Material and Communication* displays an inverted-U shape in the data, which can be captured by the non-homothetic CES. Finally, public transportation expenditures are roughly flat in the data or slightly increasing. My estimated ψ_j implies a slightly flatter pattern than observed in the data.

The consumption behavior with respect to more luxurious categories (with a higher expenditure elasticity), too, is captured well by the non-homothetic CES preferences. *Food away from home* constitutes the largest expenditure item in this category – in particular for high-income households which allocate more than 10% of their consumption budget to this category. *Recreation and sports*, *Personal services*, and *Domestic services* are all categories that one would intuitively identify as luxuries and the estimated consumption patterns point in a similar direction with expenditure shares increasing strongly in income. Potentially due to the salient externalities of *Alcohol and tobacco*, this category's consumption pattern in the data deviates slightly from the one resulting from the model. But given the small contribution to total expenditures, this is likely not to drive my quantitative results. Overall, I conclude that the non-homothetic CES estimated on US household data provide a very good fit of the consumption behavior observed along the income distribution.

4 Optimal Consumption Taxes

In this section, I show that consumption taxes are optimally non-uniform and accompanied by decreases in labor income tax progressivity as compared to the *status quo*. The conclusion is derived by solving a Ramsey problem in which the government optimally chooses consumption and labor income taxes so as to maximize a utilitarian welfare function.

To limit the dimensions of the optimization problem and to streamline the discussion of the optimal design of consumption taxes, I assume that there are two different consumption tax rates: one for necessities and one for luxuries. This is very much in line with what is observed empirically where most countries have 2 or at most 3 different consumption tax rates (Ebrill et al., 2001). To distinguish necessities and luxuries, I use the expenditure elasticity of the various consumption categories with respect to changes in the expenditures allocated to consumption expenditures.[17] Whenever this elasticity is smaller than 1, I consider the good or service to be a necessity good and conversely goods and services with an elasticity larger than 1 are considered luxuries. Since the non-homothetic CES preferences generate income-dependent expenditure elasticities, I use the elasticity of the median income household in the sample to determine the split between necessities and luxuries. Figure 4 plots these elasticities in ascending order.

4 of the 11 categories are considered necessities and 7 are applied the luxury rate. For the median income household, this corresponds to a split of exactly 50%-50% of necessities and luxuries.

4.1 Ramsey Problem

In this section, I show that the optimal design of consumption taxes is characterized by relatively high taxes on luxury goods and a strong subsidy on necessity goods, resulting in a smaller economy with higher labor income taxes. To do so, I solve a Ramsey problem in which the government optimizes over consumption taxes and labor income taxes.

The government plan is fully characterized by $\boldsymbol{\tau} = \{\tau_n, \tau_l, \tau, \lambda\}$; the consumption taxes on the two types of consumption goods, the progressivity of the income tax, and the level of the income tax. The emphasis of this paper is on the design of consumption taxes and their interaction with the labor income tax. Hence, I focus on the optimal τ_n and τ_l as well as the optimal progressivity of the labor income tax. The level of the labor income tax λ is used to clear the government budget constraint. For all of these exercises, the capital income tax is fixed at the calibrated value of 36%. I solve for the optimal consumption tax by taking into account the full transition from the initial

[17]Formally, this elasticity is defined as:

$$\eta_{E_j,E_{\mathcal{C}}} = \frac{\partial E_j}{\partial E_{\mathcal{C}}} \frac{E_{\mathcal{C}}}{E_j}$$

where $E_j = c_j p_j (1 + \tau_j)$.

FIGURE 4: Expenditure Elasticities

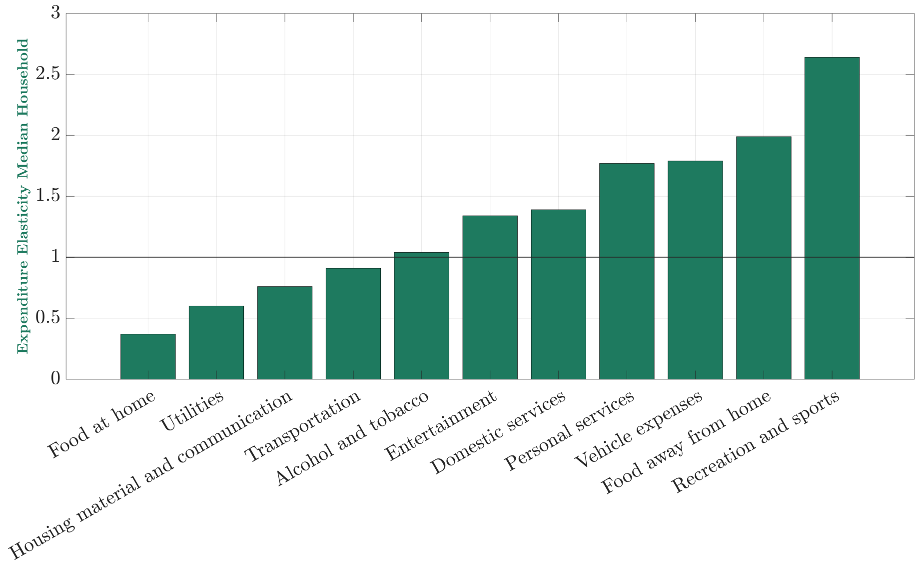

Note: This figure presents expenditure elasticities with respect to changes in E_χ of the median income household. Goods and services with elasticity < 1 are taxed at a rate $\tau_{e,n}$ and goods and services with elasticity > 1 at $\tau_{e,l}$, respectively.

steady state to the final steady state. Adjustments in the capital stock over the transition can substantially alter the welfare conclusions since it may induce higher or lower consumption during the transition in order to accumulate or diminish the capital stock (Auerbach and Kotlikoff, 1987).

The generic government problem is defined as follows:

$$\max_{\boldsymbol{\tau}} W(\boldsymbol{\tau}) = \int_{\mathcal{Z}\times\mathcal{A}\times\mathcal{S}} v_0\left(z,a,s;\boldsymbol{\tau}\right) d\mu_0\left(z,a,s\right) \tag{7}$$

The baseline results assume that the benevolent government aggregates households' individual utilities using a Utilitarian social welfare function with equal weights assigned to each household. Welfare is computed using the distribution of households at the initial steady state $\mu_0(z,a,s)$ and by aggregating life-time utility of households from the time of the policy reform to infinity.[18]

Table 4 displays the optimal policy under different constraints on which tax instruments the government can vary. If the government is constrained to set uniform consumption taxes and the labor income tax progressivity remains at its current level, consumption is subsidized with a negative tax rate of 26%. Allowing the government to vary the tax on necessities and luxuries, the optimal policy suggests that consumption taxes should be highly non-uniform. While necessities are subsidized at a rate of 40%, luxuries are exempted from the tax. The fact that consumption of the outside good is non-homothetic, with low-income households allocating a higher share of

[18]In the numerical approach to finding the optimal tax code, I follow Boar and Midrigan (2022) and evaluate the welfare function on a coarse grid of the tax parameters. Using the best points from that grid, I then conduct a local search around that point.

TABLE 4: Optimal Policies

	Tax Instrument			
	τ_n	τ_l	τ	λ
Benchmark	0%	5%	0.12	0.95
$\tau_n = \tau_l, \tau = \tau_{SQ}$	-26%	-26%	0.12	0.63
$\tau = \tau_{SQ}$	-40%	0%	0.12	0.55
Optimal	-52%	7%	-0.07	0.37

Note: This table summarizes the results originating from solving the problem of the benevolent government (Equation 7). *Benchmark* corresponds the policy under the calibrated steady state. The second row is the optimal policy when there is only one uniform consumption tax and the labor income tax progressivity is constrained to the *status quo*. The third row allows for differentiated consumption taxes, and the *Optimal* row is the final optimal policy where all tax instruments can be set freely.

their expenditures to consumption goods, is likely to play a crucial role in this result. High-income households do spend less on luxurious goods, which results in a higher rate for luxuries than necessities, but given the asymmetric allocation across outside goods and consumption goods, the rate on luxuries cannot be too high. This avoids applying a too high average consumption tax rate to low-income households, which also consume *some* luxurious goods. Allowing the planner to vary both consumption taxes and the progressivity of the labor income tax, the optimal policy prescribes a decrease in the progressivity[19] of the labor income tax and highly non-uniform consumption tax rates of 7% for luxury goods and -52% for necessities. Since the subsidy for necessities requires additional government revenue, the level of the labor income tax is increased markedly, which is mirrored in a reduction of λ from 0.95 to 0.37. Overall, the policy provides a welfare gain of 1.6% as measured in consumption equivalent variation (CEV).[20] Next, I will analyze various features of this optimal policy to get an understanding of what drives this large rate differentiation.

Aggregate Dynamics. Figure 5 summarizes the transitional dynamics of some key aggregate variables and some measures related to consumption inequality. In terms of aggregate variables, this reform to the optimal rates results in a substantial decrease of the capital stock by 8%, lower aggregate labor supply (aggregate effective hours reduce by 2.7%) and with this a smaller total size of the economy, which shrinks by around 4.1%. These changes in aggregate numbers also results in a slight reduction of wages and an increase in the return, thereby shifting income in favor of capital holders. As mentioned earlier, the subsidy on necessities must be matched by an increase in the labor income tax, as shown in the first plot of the second row. While there is an initial jump

[19] Despite the fact that τ turns out to be negative, the presence of the lump-sum transfer keeps the tax schedule progressive.

[20] Note that I consider CEV as the percentage increase in total consumption (so both the change in each c_j as well as in c_o) that a household would have to been given in the initial steady state such that she is indifferent between the initial steady state and the policy reform including the entire reform. Positive values therefore indicate that the household benefits from the reform.

FIGURE 5: Transitional Dynamics

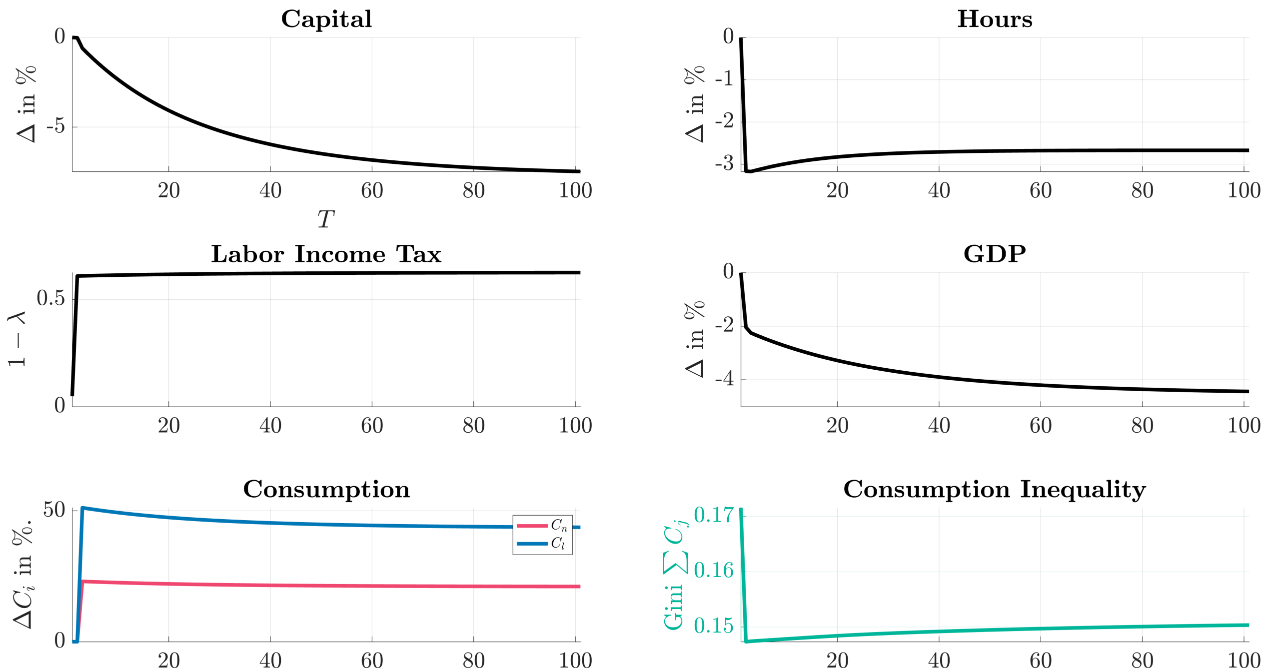

Note: Transitional dynamics of aggregate variables from initial steady state to the optimal tax system. All plots except the bottom row show the % deviation from the initial steady state.

at the time of the implementation of the policy, the level then slowly continues to increase until it reaches its new steady state value.

In terms of consumption, the policy reform results in an increase of both the aggregate consumption of necessities and luxuries. Interestingly, despite the large subsidy on necessities and the tax on luxuries, the economy shifts towards an aggregate consumption basket that features a higher share of luxury goods and services as compared to the initial steady state. The reason is that the subsidy on necessities frees up disposable income, which can be used to finance more luxury consumption. Because the demand for necessities is less elastic than the demand for luxuries, we observe a strong increase of the aggregate consumption of luxuries by a bit less than 50% as compared to an increase in aggregate necessities consumption of only slightly above 20%. Finally, the policy decreases substantially consumption inequality as measured by the Gini index of consumption of pure consumption goods.[21] Of course, the increase in consumption of goods and services, combined with the reduction in GDP, means that households shift their expenditures from the outside good to consumption goods.

Partial Reforms. Table 5 summarizes the change in various aggregate numbers when reforming only one of the tax parameters. I proceed by subsequently reforming each of the tax instruments. First, I change the tax on necessities to the optimal level of -52%. Next, I also adjust the tax on

[21]Note that the same pattern is observed if the consumption of the outside good is included in the computation of the Gini index.

TABLE 5: Partial Reforms

	Partial Reforms			
Variable	Benchmark	$+\Delta\tau_n$	$+\Delta\tau_l$	$+\Delta\tau$
Taxes				
τ_n	0.0%	-52%	-52%	-52%
τ_l	5%	5%	7%	7%
τ	0.12	0.12	0.12	-0.07
λ	0.95	0.53	0.54	0.37
Aggregate				
Capital	103.5	-22.6%	+0.56%	+18.7%
Hours	1.9	-8.6%	+0.2%	+6.3%
GDP	27.36	-13.9%	+0.3%	+10.6%
Wage	9.0	-5.8%	+0.12%	+4.07%
Return	7.98%	+1.07pp.	-0.02pp.	-0.72pp.
Inequality				
Wealth Gini	0.55	-0.01	+0.00	+0.02
$Earnings\ (after-tax)$	0.39	+0.02	+0.00	+0.07
$\sum C_j\ Gini$	0.17	-0.04	-0.01	+0.02
$\sum C_j + C_o\ Gini$	0.23	-0.03	+0.00	+0.03

Note: The numbers are obtained by analyzing the change in steady values between the initial steady state and the reform, which only includes changes in one tax parameter.

luxuries from 5% to 7%. Finally, I decrease the progressivity of the labor income tax from 0.12 to -0.07.

Overall, the table suggests that the largest changes in aggregate numbers derive from the large subsidy introduced on necessity goods but also the change to the labor income tax. The change in the tax on necessities in isolation reduces GDP by more than 13%, increases the level of the labor income tax, and reduces consumption inequality. Compared to this, the increase in the tax on luxuries and the decrease in labor income tax progressivity have smaller effects on aggregate numbers. For instance, increasing the tax on luxuries mildly increases the total number of hours worked and incentivizes slightly capital accumulation. Conversely, the decrease in labor income tax progressivity counteracts the initial effects of the subsidy on necessities by stimulating capital accumulation and labor supply. This pattern is also reflected in the factor prices, where we observe a large increase of more than 1 percentage point in the return to capital from changing the tax on necessities and a reduction of the wage by more than 5% in response to the introduction of the subsidy on necessities. In contrast, changing the tax on luxuries and labor income tax progressivity partially offset these effects. The table also highlights that the policy reform is successful in decreasing consumption inequality but that this comes at the cost of increased earnings inequality. This highlights exactly the shift from providing insurance to labor income to insuring consumption instead.

Distributional Effects. The changes in factor prices, the large rate differentiation, and the change to the labor income tax are likely to have differential effects on households depending on their wealth, skill type, and idiosyncratic productivity. To better understand these differential effects, Figure 6 plots the change in welfare for different types of households as measured by consumption equivalent variation.

FIGURE 6: Welfare Changes

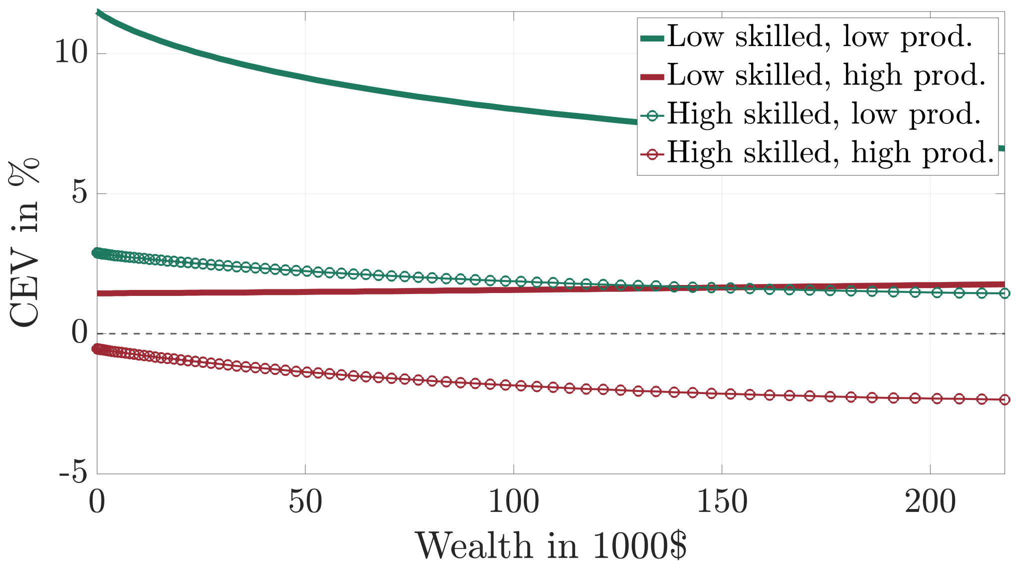

Note: Changes in welfare (measured in Consumption Equivalent Variation) as a function of household wealth, skill type and idiosyncratic productivity. Welfare change computed when changing from the initial calibrated steady state to the optimal policy.

The large gains in welfare are concentrated in some particular parts of the population. It is, in particular, the unproductive households with little wealth who benefit extensively from the reform. Households with low labor productivity and little wealth consume consumption baskets dominated by necessities. Under the optimal policy, they receive large subsidies on a large share of their consumption. This allows them to substantially increase their consumption of both necessities and luxuries. Given that they have little income overall, the increase in the labor income tax hurts them only to a minor extent. The welfare gains are clearly decreasing in household wealth. As will be shown below, the share of consumption going to luxuries is increasing in wealth and the higher tax on luxuries therefore adversely affects households with higher wealth position.

The only households that clearly lose from the policy reform are very productive households that are of the high-skill type. Those households' consumption baskets contain large amounts of luxuries, and the increase in the tax on luxuries, the higher level of the labor income tax, and the reduction in wages therefore affect these households through three channels, resulting in an overall loss in welfare from the policy reform. Nonetheless, the reform generates an overall welfare gain of more than 1.6% and more than 60% of the population benefits from the reform. In other words, if the population were to vote by simple majority vote on whether such a reform should be implemented, the result would be clearly in favor of the reform.

5 Mechanisms

Having analyzed the aggregate and distributional consequences of switching to the optimal policy, I now zoom into three distinct channels resulting in the optimal policy of $\tau_n = -52\%$, $\tau_l = 7\%$ and $\tau = -0.07$: redistribution, labor supply effects, and consumption taxes as an implicit tax on initial wealth.

Redistribution and Insurance. A large argument in favor of non-uniform consumption taxes stems from redistribution and consumption insurance. On the one hand, reducing the average consumption tax of low-income households redistributes resources from households with low marginal utilities to households with high marginal utilities, thereby resulting in a net benefit in terms of aggregate welfare. On the other hand, providing a subsidy on necessity goods acts as consumption insurance since goods and services primarily consumed by households that experience an adverse idiosyncratic shock are taxed at a low (or even negative) rate.

FIGURE 7: Consumption Changes

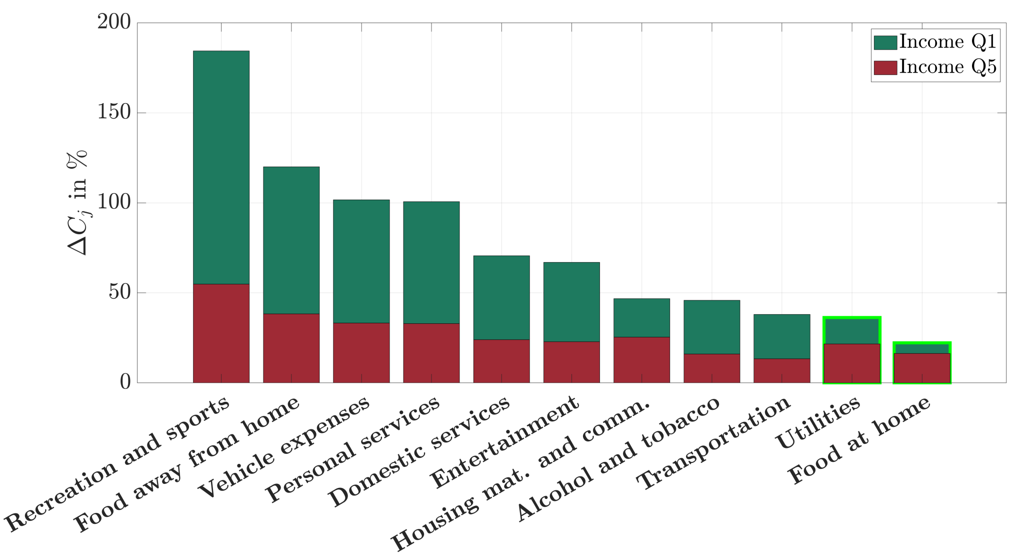

Note: Changes in consumption of the 11 consumption categories of households in the 1^{st} and 5^{th} income quintile in response to switching from the calibrated initial steady state to the optimal policy. Bars with a light-green outer line are necessities whereas the other categories are applied the tax for luxuries.

Figure 7 highlights this channel by plotting the percentage change in quantities consumed of each of the 11 consumption categories by households in the first and 5^{th} income quintile. It is clear that both types of households increase substantially their consumption levels in response to the switch to the optimal tax policy. Changes in the distribution of available income affect the overall share of expenditures that go to consumption and to the outside good. This is due to the fact that the split between these two types of expenditures is non-homothetic in itself. Hence, the policy reform induces an increase in consumption of pure consumption goods even though aggregate GDP decreases. This is due to the fact that some households switch their consumption from consumption goods to the outside good.

The increase is, however, distributed unevenly both between households and across goods. Overall, the reform redistributes towards low-income households and it is for this reason that they are able to increase their consumption substantially more than households in the 5^{th} income quintile. In line with what has been found when analyzing the aggregate effects of the reform, the figure highlights that it is in particular the luxury goods (bars without a light-green frame), whose consumption is expanded significantly. The subsidy on necessity goods frees up resources for these households and allows them to increase the consumption of the more elastic luxury goods, such as *Recreation and Sports* or *Food away from home*.

Labor Supply Effects. The interaction between labor income taxes and differentiated consumption taxes has been studied extensively. For instance, Atkinson and Stiglitz (1976) argue that the latter should not be implemented and, instead, all redistribution should take place through a non-linear labor income tax. One might therefore wonder why the optimal policy is characterized by only a minor increase in labor income tax progressivity and very non-uniform consumption taxes. First, in this model, there are more sources of income than just labor income, and non-uniform consumption taxes can go after those alternative sources of income. The next section will discuss this in more detail. Second, in this model, changing the progressivity of consumption taxes by introducing non-uniform consumption taxes is not equivalent to changing the progressivity of the labor income tax.

It is useful to consider the optimality condition for labor:

$$n^{\chi} = U'(E)(1 - \tilde{\tau}_l)wz \tag{8}$$

where without loss of generality I set $B_l = 1$ and defined $\tilde{\tau}_l \equiv 1 - (1-\tau)\lambda(wzsn)^{-\tau}$ as the marginal labor income tax rate. Consider the optimal labor supply in the homothetic case with one good, no outside expenditures, and the price of the outside good normalized to 1. In that case, the marginal utility of consumption (or expenditures) is given by $U'(E) = \frac{c^{-\theta}}{1+\tau_c}$ and following Boar and Midrigan (2022) we can implicitly define the labor wedge as: $\nu_l \equiv \frac{(1-\tilde{\tau}_l)}{1+\tau_c}$. Increasing the progressivity of the labor income tax (increasing τ) disincentivizes labor supply of households by increasing the labor wedge. Similarly, increasing the the consumption tax affects the denominator of the labor wedge. In the fully-fledged non-homothetic specification, the marginal utility from total expenditures takes the more involved form:

$$U'(E) = \left[\mathcal{C}^{\eta}(c_o - \underline{c}_o)^{1-\eta}\right]^{-\theta} \frac{\partial}{\partial E}(\mathcal{C}^{\eta}(c_o - \underline{c}_o)^{1-\eta})$$

Here, the marginal utility with respect to total expenditures depends on how the split between outside and consumption expenditures responds to changes in expenditures, but – of course – also how the split between different consumption goods changes. Even more so, the marginal utility hinges on the current consumption basket, which – in turn – affects labor supply responses. Figure 8 presents simulated labor supply responses to illustrate this point. The first subplot shows the response in hours worked when increasing the progressivity of the labor income tax. The second and third plots, respectively, show labor supply responses to increases in either τ_n or τ_l. All these responses are in partial equilibrium and do not take into account changes in factor prices.

FIGURE 8: Labor Supply Changes for $\Delta\tau$, $\Delta\tau_n$, and $\Delta\tau_l$

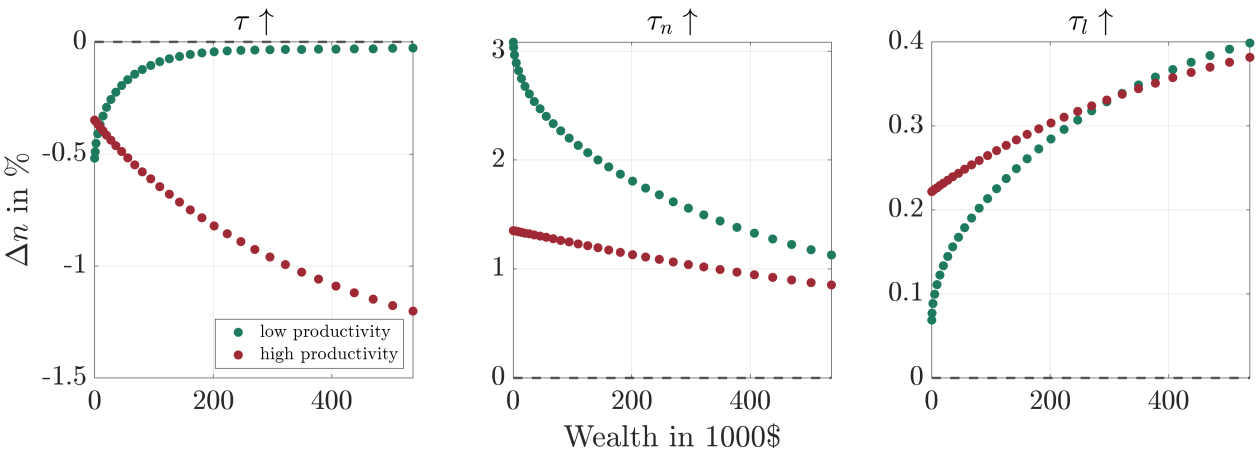

Note: Changes in labor supply of two types of households in response to increasing the progressivity of the labor income tax (left plot), the level of the tax on necessities (middle plot), and the tax on luxuries (right plot). All responses are in partial equilibrium without changing factor prices or any other tax instruments.

The simulation indicates that increasing the differential between the tax on luxuries and necessities and changing the progressivity of the labor income tax are distinct. Looking at the differential response of households with different levels of productivity, one can observe that it is in particular the high-productivity households, which decrease their labor supply in response to increases in τ. Clearly, the decrease in the progressivity of the labor income tax therefore incentivizes highly productive households to increase their labor supply.

As can be seen in the two subsequent plots of Figure 8, increasing the tax on necessities or luxuries has very different effects than increasing the progressivity of the labor income tax. In particular, in comparison to very productive households, low-productivity households – whose consumption basket is dominated by necessities – increase their labor supply more in response to hikes in the tax on necessities. Conversely, very productive households respond more when the tax on luxuries is increased. This highlights that decreasing the tax on necessities and increasing the tax on luxuries has differential effects on the labor supply of different households. The large subsidy on necessities incentivizes all households to decrease their labor supply, but the effect is particularly strong for unproductive households. Conversely, the increase in the tax on luxuries triggers an increase in hours worked of very productive households. Finally, the reduction in τ adds to this by incentivizing labor supply of high-producitivty households. These three changes combined, therefore, shift the labor supply towards the very productive households, thereby resulting in an overall efficiency gain.

Figure 9 summarizes this argument in one plot by plotting the differences in labor supply of dif-

FIGURE 9: Change in Hours Worked (General Equilibrium)

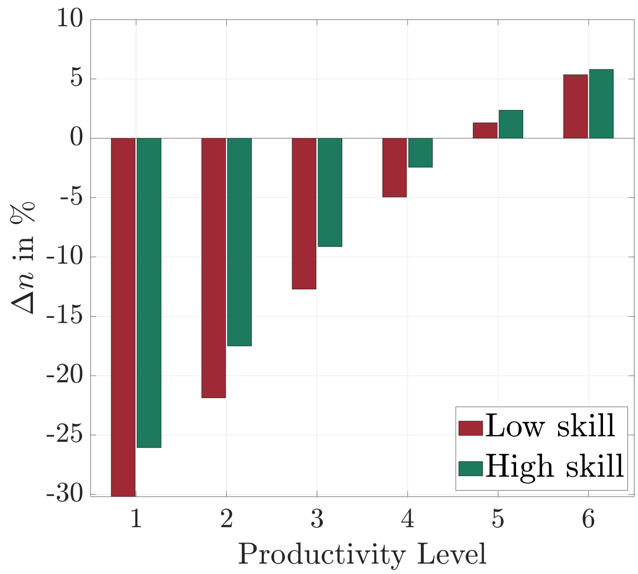

Note: Change in hours worked by skill type and productivity level when moving from the calibrated steady state to the optimal policy. As compared to Figure 8, this plot takes into account all general equilibrium changes in factor prices.

ferent households. These results represent general equilibrium responses, reflecting all factor price adjustments between the initial and final steady states. The policy experiment entails adjusting the three tax instruments from their baseline calibration to their respective optimal levels. In response to this reform, labor supply reallocates toward highly productive households. Less productive households reduce their labor supply, as the relative affordability of their primary consumption good increases. By contrast, highly productive households face a higher tax burden on the goods and services they most prefer, but also a labor income tax schedule that has become substantially less progressive. Taken together, these changes incentivize productive households to increase their labor supply by extending their working hours. This mechanism illustrates the efficiency gains generated by the reform.

Implicit Tax on Initial Wealth. In the optimal tax problem, the government can only adjust consumption taxes and labor income taxes. Nonetheless, the planner can implicitly impose a tax on high-wealth households by taxing at a higher rate goods and services favored by wealthy households. Since Nishiyama and Smetters (2005), it is understood that a consumption tax increase operates as an implicit tax on pre-existing wealth when the economy transitions from a steady state without a consumption tax to one with a positive rate. Agents holding wealth at the time of implementation have no alternative but to consume it and thereby incur the tax. Because adjustment margins are limited at the time of the policy change, the induced distortions are relatively small. With non-homothetic preferences and heterogeneous consumption baskets, the same logic applies, and the implicit tax on initial wealth can even be directed disproportionately toward high-wealth households. Figure 6 illustrates this mechanism using two representative households.

TABLE 6: Consumption of Two Example Households

	Households	
	low wealth	high wealth
Labor income	1.55	1.55
Wealth	0.87	234.69
Post-tax income	5.49	17.54
Capital income	0.45	12.15
Consumption share necessities	84%	60%
Consumption share share luxuries	16%	39%

Note: This table presents various characteristics of two example households. They have identical levels of labor income but very different wealth positions.

The figure illustrates two households in the model economy with identical labor income but markedly different wealth positions. The household with substantial wealth receives additional capital income and, as a result, allocates a smaller share of expenditure to necessities (60%) compared to the household with little wealth (84%). The former devotes a larger fraction of its budget to categories such as *Recreation and sports* or *Entertainment*, whereas the latter allocates most to *Food at home* or *Utilities*. Differential tax rates across consumption goods, therefore, imply an implicit wealth tax that falls more heavily on high-wealth households. Whereas a labor income tax targets households with high labor earnings, a differentiated consumption tax effectively targets households with high wealth, while generating only limited distortions. Although a uniform consumption tax would also impose a higher absolute burden on high-wealth households, differentiation allows for mitigation of adverse effects on low-wealth, low-income households. The policy thus combines efficiency – owing to the limited scope for behavioral responses – with equity, by shifting the tax incidence toward households with large wealth holdings.

6 CONCLUSION

This paper analyzes the optimal joint design of consumption and labor income taxes in a quantitative heterogeneous-agent model with non-homothetic preferences. The model is estimated to match various features of the US economy. Solving a Ramsey problem in which the planner optimizes over both consumption tax rates and income taxes yields the central result: optimal consumption taxes are markedly non-uniform. Necessities are subsidized while luxuries are taxed at positive rates. This outcome is driven by three mechanisms: subsidized necessities provide insurance to households exposed to adverse idiosyncratic shocks, taxation of luxuries serves as an implicit tax on previously accumulated wealth, and differentiated rates strengthen labor supply incentives among highly productive households.

The findings contribute to the literature on optimal taxation in two respects. First, they demonstrate that even in the presence of an optimally designed labor income tax, differentiated consumption taxes can generate sizable welfare gains. This underscores that consumption and income taxes are not perfect substitutes once endogenously arising wealth heterogeneity and non-homothetic preferences are taken into account. Second, they highlight the redistributive potential of commodity taxation in environments where wealth taxation is either politically or administratively infeasible.

More broadly, the analysis suggests that consumption should not be treated as a homogeneous aggregate when evaluating tax policy. The heterogeneity of consumption baskets across the income and wealth distribution provides a channel through which differentiated consumption taxes can improve both equity and efficiency. These insights help rationalize the widespread use of differentiated rates in practice and point to their potential role in modern tax systems where traditional redistributive instruments may be constrained.

References

Aguiar, Mark and Mark Bils, "Has Consumption Inequality Mirrored Income Inequality?," *American Economic Review*, 2015, *105* (9), 2725–2756.

Aiyagari, S Rao, "Uninsured Idiosyncratic Risk and Aggregate Saving," *The Quarterly Journal of Economics*, 1994, *109* (3), 659–684.

_ , "Optimal Capital Income Taxation with Incomplete Markets, Borrowing Constraints, and Constant Discounting," *Journal of Political Economy*, 1995, *103* (6), 1158–1175.

Atkinson, Anthony B. and Joseph E. Stiglitz, "The design of tax structure: Direct versus indirect taxation," *Journal of Public Economics*, 1976, *6* (1-2), 55–75.

Auerbach, Alan J. and Laurence J. Kotlikoff, "Evaluating Fiscal Policy with a Dynamic Simulation Model," *American Economic Review*, 1987, *77* (2), 49–55.

Bastani, Spencer and Sebastian Köhne, "How Should Consumption Be Taxed?," *Ifo Working Paper*, 2022, *10038.*

Bewley, Truman, "A Difficulty with the Optimum Quantity of Money," *Econometrica*, 1983, *51* (5), 1485.

Bhattarai, Saroj, Woong Yong Park, Jae Won Lee, and Choongryul Yang, "Macroeconomic Effects of Capital Tax Rate Changes," *Fed Finance and Economics Discussion Series*, 2022, *2022-027.*

Boar, Corina and Virgiliu Midrigan, "Efficient Redistribution," *Journal of Monetary Economics*, 2022, *131*, 78–91.

Borusyak, Kirill and Xavier Jaravel, "The Distributional Effects of Trade: Theory and Evidence from the United States," *NBER Working Paper*, 2022, *28957.*

Bureau of Labor Statistics, "Consumer Expenditure Survey, 2001–2019," https://www.bls.gov/cex/ 2020. Accessed: 2025-08-19.

Bénabou, Roland, "Tax and Education Policy in a Heterogeneous-Agent Economy: What Levels of Redistribution Maximize Growth and Efficiency?," *Econometrica*, 2002, *70* (2), 481–517.

Carroll, Christopher D., "The method of endogenous gridpoints for solving dynamic stochastic optimization problems," *Economics Letters*, 2006, *91* (3), 312–320.

Carroll, Daniel, André Victor Doherty Luduvice, and Eric R. Young, "Optimal Fiscal Reform with Many Taxes," *Federal Reserve Bank of Cleveland Working Paper Series*, 2024, *Working Paper No. 23-07R.*

Chari, V.V. and Patrick J. Kehoe, "Chapter 26 Optimal fiscal and monetary policy," *Handbook of Macroeconomics*, 1999, *1*, 1671–1745.

Comin, Diego, Ana Danieli, and Marti Mestieri, "Income-Driven Labor Market Polarization," *Federal Reserve Bank of Chicago Working Paper Series*, 2020, *WP 2020-22.*

_ , Danial Lashkari, and Martí Mestieri, "Structural Change With Long-Run Income and Price Effects," *Econometrica*, 2021, *89* (1), 311–374.

Conesa, Juan Carlos and Dirk Krueger, "On the optimal progressivity of the income tax code," *Journal of Monetary Economics*, 2006, *53* (7), 1425–1450.

_ , Bo Li, and Qian Li, "Welfare Implications of Switching to Consumption Taxation," *Journal of Economic Dynamics and Control*, 2020, *120*, 103991.

Congressional Budget Office, "The Distribution of Household Income, 2019," https://www.cbo.gov/system/files/2022-11/58353-HouseholdIncome.pdf 2022.

Diamond, Peter A. and James A. Mirrlees, "Optimal Taxation and Public Production I: Production Efficiency," *American Economic Review*, 1971, *61* (1), 8–27.

Domeij, David and Jonathan Heathcote, "On The Distributional Effects Of Reducing Capital Taxes," *International Economic Review*, 2004, *45* (2), 523–554.

Dyrda, Sebastian and Marcelo Pedroni, "Optimal Fiscal Policy in a Model with Uninsurable Idiosyncratic Shocks," *The Review of Economic Studies*, 2022, *90* (2), 744–780.

Ebrill, Liam, Michael Keen, and Victoria Summers, *The Modern VAT*, International Monetary Fund, 2001. ch 7 has stuff on rate differentition, esp see box 7.2.

Farmer, Leland E. and Alexis Akira Toda, "Discretizing nonlinear, non-Gaussian Markov processes with exact conditional moments," *Quantitative Economics*, 2017, *8* (2), 651–683.

Federal Reserve Board, "Survey of Consumer Finances, 2013," https://www.federalreserve.gov/econres/scfindex.htm 2013. Accessed: 2025-08-19.

Ferriere, Axelle, Philipp Grübener, Gaston Navarro, and Oliko Vardishvili, "On the Optimal Design of Transfers and Income Tax Progressivity," *Journal of Political Economy Macroeconomics*, 2023, *1* (2).

FRED, "Is college still worth it?," https://fredblog.stlouisfed.org/2018/07/is-college-still-worth-it/ 2018.

Golosov, Mikhail, Maxim Troshkin, Aleh Tsyvinski, and Matthew Weinzierl, "Preference heterogeneity and optimal capital income taxation," *Journal of Public Economics*, 2013, *97*, 160–175.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L Violante, "Optimal Tax Progressivity: An Analytical Framework," *The Quarterly Journal of Economics*, 2017, *132* (4), 1693–1754.

Huggett, Mark, "The Risk-free Rate in Heterogeneous-agent Incomplete-insurance Economies," *Journal of Economic Dynamics and Control*, 1993, *17* (5-6), 953–969.

Kindermann, Fabian and Dirk Krueger, "High Marginal Tax Rates on the Top 1 Percent? Lessons from a Life-Cycle Model with Idiosyncratic Income Risk," *American Economic Journal: Macroeconomics*, 2022, *14* (2), 319–366.

Krusell, Per, Vincenzo Quadrini, and José-Víctor Ríos-Rull, "Are consumption taxes really better than income taxes?," *Journal of Monetary Economics*, 1996, *37* (3), 475–503.

Macnamara, Patrick, Myroslav Pidkuyko, and Raffaele Rossi, "Taxing Consumption in Unequal Economies," *University of Manchester Economics Discussion Paper Series*, 2022, *EDP-2210.*

Nardi, Mariacristina De, Giulio Fella, and Gonzalo Paz-Pardo, "Nonlinear Household Earnings Dynamics, Self-Insurance, and Welfare," *Journal of the European Economic Association*, 2019, *18* (2), 890–926.

Nishiyama, Shinichi and Kent Smetters, "Consumption Taxes and Economic Efficiency with Idiosyncratic Wage Shocks," *Journal of Political Economy*, 2005, *113* (5), 1088–1115.

Nord, Lukas, Marcel Peruffo, and Caterina Mendicino, "Distributive Effects of Banking Sector Losses," *SSRN Electronic Journal*, 2025.

Parodi, Francesca, "Taxation of Consumption and Labor Income: a Quantitative Approach," *American Economic Journal: Macroeconomics*, 2023, *15* (4), 177–216.

Ramsey, Frank, "A Contribution to the Theory of Taxation," *The Economic Journal*, 1927, *37* (145), 47.

Saez, Emmanuel, "The Desirability of Commodity Taxation Under Non-linear Income Taxation and Heterogeneous Tastes," *Journal of Public Economics*, 2002, *83* (2), 217–230.

Trabandt, Mathias and Harald Uhlig, "The Laffer curve revisited," *Journal of Monetary Economics*, 2011, *58* (4), 305–327.

US Census Bureau, "CPS Historical Time Series Tables," https://www.census.gov/data/tables/time-series/demo/educational-attainment/cps-historical-time-series.html 2023.

Vardishvili, Oliko, "The Macroeconomic Cost of College Dropouts," 2020.

Young, Eric R., "Solving the incomplete markets model with aggregate uncertainty using the Krusell–Smith algorithm and non-stochastic simulations," *Journal of Economic Dynamics and Control*, 2010, *34* (1), 36–41.

Özlem Kina, Ctirad Slavik, and Hakki Yazici, "Redistributive Capital Taxation Revisited," *American Economics Journal: Macroeconomics*, 2024, *16* (2), 182–216.

A Appendix

A.1 Equilibrium Definition

I denote as $\mathcal{A}$ the space for assets, $\mathcal{Z}$ the space for households productivity and $\mathcal{S}$ the space for skill-types in the economy. The state space is then defined as $Q = \mathcal{A} \times \mathcal{Z} \times \mathcal{S}$, and I let $\mathcal{B}$ be the Borel σ-algebra induced by Q.

Definition Competitive Equilibrium. A competitive equilibrium consists of: (i) aggregate prices w_t, r_t, $\{p_{j,t}\}_{j\in J}$, $p_{o,t}$, (ii) labor supply, savings, and consumption decisions for the J goods and the outside good $n_t(a,z,s)$, $a_{t+1}(a,z,s)$, $\{c_{j,t}(a,z,s)\}_{j\in J}$, $c_{o,t}(a,z,s)$ and the value function $V_t(a,z,s)$, (iii) employment, capital and output choices of the basic goods firm L_t, K_t, $Y_{0,t}$, (iv) measures of households over their states $\mu_t(a,z,s)$ over $\mathcal{B}$, and (v) government decisions $\{G_t, D_t, \mathcal{T}_t(\cdot), \tau_{k,t}, \tau_{o,t}, \{\tau_{j,t}\}_{j\in J}\}$ such that:

1. Household policies result from solving the household problem given aggregate prices.
2. Firm choices result from solving the firm problem given aggregate prices.
3. The government budget constraints holds period-by-period.
4. All markets clear.
   - Labor market clearing condition:
   $$L_t = \int_Q zsn_t(a,z,s)d\mu_t(a,z,s)$$
   - Asset market clearning condition
   $$K_{t+1} = \int_Q a_{t+1}(a,z,s)d\mu_t(a,z,s)$$
   - Goods market clears through Walras' Law.
5. The measure μ_t is consistent with households' policies in the sense that $\mu_{t+1}(\mathcal{B}) = \mathbb{1}_{\{a_{t+1}(a,z,s)\in\mathcal{B}\}} \sum_{z_{t+1}\in\mathcal{Z}} \Pi(z_{t+1}|z_t)$

A.2 Computational Details

Steady State

To solve for the steady state of the economy, I go through the following steps:

1. Provide a guess for aggregate capital K, aggregate labor L and the government budget constraint clearing level of labor income taxes λ.
2. Compute factor prices w and r using the intermediate good firm's first order condition.

3. Solve the household problem. I use a modified version of the Endogenous Grid Method (Carroll, 2006) to solve the household problem. In particular, I iteratively use the Euler equation and the labor optimality condition to solve for optimal policies for savings, expenditures, and labor supply.

4. Aggregate households using the Young (2010) method. Obtain aggregate labor supply L^s and aggregate capital supply K^s. Solving the government budget constraint for the level of labor income taxes, which would make the constraint hold with equality, I can also derive a λ^s.

5. Check whether $K \approx K^s$, $L \approx L^s$, and $\lambda \approx \lambda^s$.

    - If yes: we have found an equilibrium.
    - If no: update the new K according to $K^{new} = \omega K + (1-\lambda)K^s$ and do the same for L and λ. To ensure convergence, ω needs to be relatively close to 1. Then, set $K = K^{new}$, $L = L^{new}$ and $\lambda = \lambda^{new}$ and go back to point (2).

Transition.
To compute the transition dynamics, I use a standard shooting algorithm. This involves first guessing a path for the equilibrium objects (K, L, λ) from the initial steady state to the final steady state. In a second step, I solve backwards from the final steady state for the policy functions and forward from the initial steady state to the final steady state the distribution $\mu(\cdot)$. I then compute implied (K, L, λ) over the transition (as in point 4-5 in the steady state computation) and compare the implied values to the given values. If they are close, I stop. Else I update using a weighted sum and recompute the transition.

A.3 Estimation Details

TABLE 7: Estimated $\{\Upsilon_j\}_{j \in J}$

Variable	Estimated Υ_j
Food at home	2.22
Utilities	1.14
Housing Material & Communication	0.51
Public Transportation	0.15
Alcohol & Tobacco	0.26
Entertainment	0.55
Domestic Services	0.59
Personal Services	0.72
Vehicle Expenses	0.51
Food away from home	1.34
Recreation & Sports	1.17

Note: This table presents the estimated share parameters ($\{\Upsilon_j\}_{j \in J}$) of the non-homothetic CES function. The parameters are estimated by targeting economy-wide, average consumption shares of the different consumption categories.

FIGURE 10: E_o and E_χ shares in the Data and in the Model

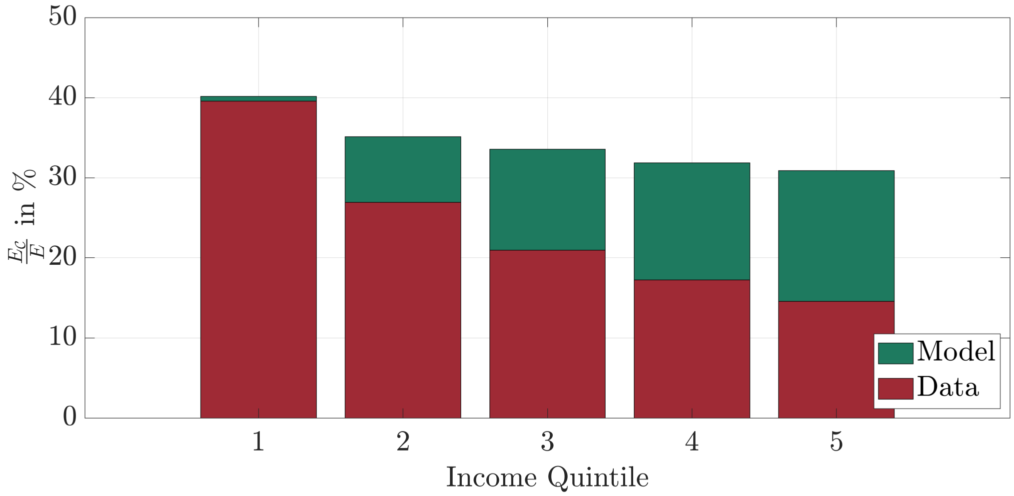

Note: This plot presents the expenditure shares allocated to E_o and E_χ. Data comes from the CEX, 2001-2019 (Bureau of Labor Statistics, 2020).

A.4 Inequality: Additional statistics

TABLE 8: Additional Inequality Measures in the Data and the Model

Earnings	Gini	P90/P10	P90/P50	P50/P10
Data	0.4	5.2	2.4	2.2
Model	0.32	5.08	1.82	2.79
Consumption Expenditures	Gini	P90/P10	P90/P50	P50/P10
Data	0.26	2.26	1.34	1.69
Model	0.19	2.38	1.53	1.55

Note: Values for earnings inequality from Vardishvili (2020), consumption inequality from CEX 2001-2019 (Bureau of Labor Statistics, 2020).

A.5 Additional Results: Policy Optimization

TABLE 9: Steady State Comparison of Expenditure Numbers

Variable	Steady State		
	Status quo	Optimal	Δ
Gini ($E_{\mathcal{C}}$)	0.19	0.20	+ 3.04 %
$P90/P10$ ($E_{\mathcal{C}}$)	2.49	2.62	+5.21%
$P90/P50$ ($E_{\mathcal{C}}$)	1.54	1.58	+2.45%
$P50/P10$ ($E_{\mathcal{C}}$)	1.61	1.65	+2.69 %
Gini (E_o)	0.26	0.30	+14.56%
$P90/P10$ (E_o)	3.71	5.01	+35.31%
$P90/P50$ (E_o)	1.75	1.91	+8.85%
$P50/P10$ (E_0)	2.11	2.62	+24.30 %
Gini ($E_o + E_{\mathcal{C}}$)	0.24	0.27	+10.53%
$P90/P10$ ($E_o + E_{\mathcal{C}}$)	3.23	3.89	+20.53%
$P90/P50$ ($E_o + E_{\mathcal{C}}$)	1.68	1.79	+6.38%
$P50/P10$ ($E_o + E_{\mathcal{C}}$)	1.91	2.17	+13.29 %

Note: This table compares various statistics related to consumption expenditure inequality in the initial and final (optimal) steady state.