Consumption Quality and Employment Across the Wealth Distribution*

Domenico Ferraro Vytautas Valaitis

January 4, 2023

Abstract

In the United States, market hours worked are approximately flat across the wealth distribution. Accounting for this phenomenon is a standing challenge for standard heterogeneous-agent macro models. In these models, wealthier households consume more and work fewer hours. We propose a theory that generates the cross-sectional wealth-hours relation as in the data. We quantify this theory in a heterogeneous-agent incomplete-markets model with three key features: a quality choice in consumption, non-homothetic preferences, and a multi-sector production structure. We show that the model produces consumption expenditure patterns consistent with the data and realistic "quality Engel curves."

JEL Classification: E24; E21; J22.
Keywords: Labor supply; Quality choice; Heterogeneity; Aggregation; Tax policy.

*Ferraro: Department of Economics, W.P. Carey School of Business, Arizona State University, PO Box 879801, Tempe, AZ 85287-9801, United States of America (e-mail: domenico.ferraro@asu.edu); Valaitis: School of Economics, University of Surrey, 388 Stag Hill, University Campus, Guildford GU2 7XH, United Kingdom (e-mail: v.valaitis@surrey.ac.uk). First version: March 24, 2020. Earlier versions circulated under the title "Wealth and Hours." We thank the Editor, Kurt Mitman, three anonymous referees, Francesco Agostinelli, David Berger, Alex Bick, Edouard Challe, Yongsung Chang, Berthold Herrendorf, Nir Jaimovich, Xincheng Qiu, Hitoshi Tsujiyama, and Gustavo Ventura, for comments and suggestions. All errors are our own.

1. Introduction

A large and important literature studies consumption and labor supply in heterogeneous-agent incomplete-markets models in which households face uninsurable wage shocks. In these models, endogenous labor supply acts as a "self-insurance" mechanism to smooth consumption, implying that work incentives taper off steeply with wealth: wealthier households consume more and work fewer hours. This prediction is at odds with the data. In the United States, employment rates and hours worked are nearly flat or mildly increasing across the wealth distribution (see Figure 1).

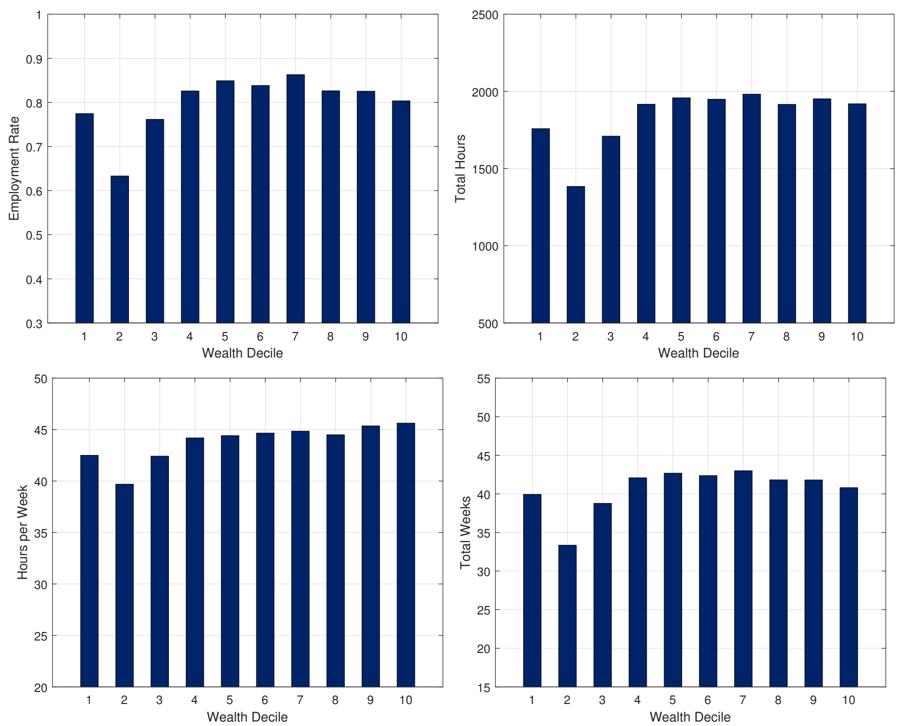

Figure 1: Employment and Hours by Wealth

Notes: The figure shows the employment rate (top left panel), total hours worked (top right panel), weekly hours worked (bottom left panel), and weeks worked (bottom right panel) by wealth deciles for household heads of 25-65 years old. Data are from the 2001-2015 Panel Study of Income Dynamics (PSID) waves. Wealth is total assets minus total liabilities at the household level. See Online Appendix A for details on variables' definitions, construction, and additional evidence.

Accounting for these cross-sectional facts is a challenge for complete- and incomplete-markets models alike. In the equilibrium of these models, wealthier households have higher consumption and, thereby, a lower marginal utility of consumption. Moreover, to the extent that leisure is a normal good, wealthier households enjoy more leisure and work fewer hours. These basic predictions hold under "balanced growth preferences" in which income and substitution effects on labor supply offset each other (King, Plosser and Rebelo, 1988), and even more so for preferences in which the income effect dominates the substitution effect.[1]

The maintained assumption in these models is that households can only choose the *quantity* of consumption, abstracting from the *quality* of the consumption basket they buy. There is, however, a growing body of empirical work suggesting that this assumption is strongly counterfactual as higher-income, wealthier households consume not only more goods but better goods (Aguiar and Bils, 2015; Bils and Klenow, 2001; Faber and Fally, 2020; Jaimovich, Rebelo and Wong, 2019; Jaimovich et al., 2019b).

This paper aims to study how "quality choice" affects labor supply and, thereby, the cross-sectional distributions of consumption expenditures, employment rates, earnings, and wealth. While quality choice has a long tradition in consumer theory (Deaton and Muellbauer, 1980), its general-equilibrium implications for labor allocations across the distribution of earnings and wealth have been unexplored. To the best of our knowledge, this is the first paper that tackles this question.

We argue that embedding quality choice into a consumption-leisure choice problem goes a long way in reconciling theory with the *cross-sectional* relationship between hours and wealth in the data. Further, the quality choice is a self-insurance mechanism in the presence of uninsurable wage shocks. In bad times, households cut back on consumption expenditures by buying cheaper lower-quality goods, with little sacrifice of the quantity of consumption. For example, think of the choice between calorie intake and the quality grades of a commodity like meat (Deaton, 1988, makes a similar analogy). We conclude that the consumption quality choice changes the standard incomplete-markets model's predictions in crucial ways. In a nutshell, abstracting from consumption quality is neither empirically plausible nor theoretically innocuous.

Our contribution to the literature is twofold. First, we propose a theory based on the quality choice in consumption that reconciles macro models with micro observations on hours and wealth. Second, we embed it in a quantitative general-equilibrium model that allows us to gauge the role of quality choice for labor allocations across the wealth distribution. In doing so, we develop a new heterogeneous-agent incomplete-markets

[1] Balanced growth preferences guarantee that household-level hours worked and aggregate hours per capita are constant along a balanced growth path. Yet, wealthier households with a larger *share* of aggregate wealth work fewer hours/exhibit lower employment rates.

model well suited for policy analysis.

Our analysis begins by studying the implications of quality choice in a static, partial-equilibrium labor supply model and a general-equilibrium model with Cobb-Douglas technology. The main idea in these models is that households choose labor supply along the extensive margin and consumption quantity and quality.[2] We derive restrictions on the class of admissible utility functions consistent with upward-sloping "quality Engel curves" that rule out homothetic preferences.[3] In addition, we show that a form of *non-separability* between the quantity and the quality of consumption is needed to preserve work incentives for the wealthy. To the extent that consumption quality rises with income and that the marginal utility of consumption depends positively on quality, wealthy households may decide to work to afford high-quality goods. In general equilibrium, however, prices of different qualities are determined to clear product markets. Ultimately, whether employment rates decrease, increase, or are constant across households with varying levels of wealth depends on their preferences for quality and the relative prices that prevail in equilibrium. In this sense, a flat relation between hours and wealth is not hard-wired into the theory.

To quantify these mechanisms, we build a heterogeneous-agent incomplete-markets model in which households face uninsurable idiosyncratic wage shocks, as in Aiyagari (1994), Bewley (1983), Huggett (1993), and Imrohoroğlu (1989). Novel features of the model are a consumption quality choice, non-homothetic preferences, and a multi-sector production structure. The incomplete-market structure allows for an endogenous wealth distribution, in which households' choices in response to idiosyncratic shocks determine their rankings in the distribution.

In the model, households face a menu of quality-price bundles from which to choose. Notably, higher-quality versions cost more. With non-homothetic preferences, the level of quality increases with income and wealth: higher-income and wealthier households consume not only more goods but also more expensive higher-quality versions (see Bils and Klenow, 2001; Faber and Fally, 2020; Jaimovich et al., 2019b, for evidence supporting this mechanism). Working long hours allows wealthy households to keep up with their preferred, high-quality consumption.

[2]We focus on the extensive margin of labor supply for two reasons. First, employment rates display the same pattern of hours worked, suggesting that the variation (or the lack thereof) in hours worked by wealth deciles is primarily accounted for by labor supply along the extensive margin. Second, there might be concerns about measurement error in self-reported hours worked; however, measurement error is arguably small if non-existent for employment rates that are measured uncontroversially.

[3]A quality Engel curve traces out unit prices against income or consumption expenditures.

The consumption quality choice also acts as a self-insurance mechanism. In the face of uninsurable, adverse wage shocks, households can cut back on quality, which allows them to maintain a relatively stable consumption quantity. This mechanism naturally interacts with labor supply decisions: wealth-poor households trade off consumption quality with the value of leisure. This mechanism is absent in the standard incomplete-markets model.

On the production side, there is a continuum of sectors producing versions of the consumption good that differ by quality and an investment sector that produces a capital good that adds to the economy's capital stock. Production functions in both industries are of the constant elasticity of substitution (CES) form. We allow factor intensities to vary by quality, encompassing the case where higher-quality goods are more labor-intensive as in Jaimovich, Rebelo and Wong (2019).

We calibrate the model to reproduce salient features of the U.S. earnings and wealth distribution and quality Engel curves estimated using Consumer Expenditure Survey (CEX) and Kilts-Nielsen data. The calibrated model yields the near-zero cross-sectional correlation between wealth and hours in the data. Also, the distribution of consumption expenditure is highly concentrated and skewed to the right, as in the data. Using data from the Panel Study of Income Dynamics (PSID), we find that this prediction lines up well with the empirical distributions of expenditure on food away from home, education, clothing, vacation, and entertainment – five good categories that are typically associated with "luxuries" (Aguiar and Bils, 2015; Chang, Hornstein and Karabarbounis, 2020).

As further validation, we use the Continuing Survey of Food Intake of Individuals (CSFII) by the U.S. Department of Agriculture. We estimate that measures of quality of food consumption based on food content (vitamins A, C, and E, calcium, cholesterol, and saturated and unsaturated fats) are greatly sensitive to household income. For example, higher-income households consume less saturated fats and cholesterol, which is typically associated with healthier and higher-quality food consumption. In contrast, the quantity of food consumption measured by total calories is virtually invariant to income.

To further highlight the role of the quality choice, we use the model to evaluate the implications for taxes and transfers. We find that quality choice with non-homothetic preferences significantly changes the standard incomplete-markets model's predictions. Overall, employment rates are more sensitive to changes in work incentives relative to the standard incomplete-market model. The reason is a much-attenuated wealth effect on labor supply, which implies less offsetting of the substitution effect. Further, a sizable share of the consumption expenditure response to changes in taxes and transfers comes from consumption quality, which interacts with the production side of the economy to the extent that higher-quality goods are more or less intensive in capital or labor.

Our paper is organized as follows. Section 2 highlights the paper's contribution to the literature. Section 3 presents theoretical results providing intuition into the consumption quality choice's implications for labor allocations. Section 4 presents the heterogeneous-agent incomplete-markets model with the consumption quality choice. We discuss the model's parameterization and quantitative properties in Sections 5 and 6, respectively. Section 7 provides evidence on consumer quality choices based on micro data. Section 8 presents two applications of the model related to taxes and transfers. Section 9 concludes.

2. Related Literature

This paper contributes to understanding labor supply under uninsurable idiosyncratic wage risk (Chang and Kim, 2006, 2007; Chang et al., 2019; Heathcote, Storesletten and Violante, 2008, 2014; Pijoan-Mas, 2006). When insurance markets are incomplete and wage shocks are persistent, labor supply acts as a self-insurance mechanism to smooth consumption. In these models, low-productivity individuals turn out to be wealth-poor. However, despite low productivity, they work as their marginal utility of consumption is high. At the same time, high-productivity individuals are wealth-rich; they do not work nearly as many hours as the wealth-poor because their marginal utility of consumption is low. This mechanism implies that work incentives taper off with wealth, generating a negative counterfactual relationship between wealth and hours.

Mustre-del-Río (2015) shows preference heterogeneity for leisure brings the standard incomplete-markets model closer to the data. In his model, wealthier households have weaker preferences for leisure. Yum (2018) argues that means-tested transfers and capital taxation play a role in reconciling the model with the data. Transfers mitigate the self-insurance motive of labor supply for wealth-poor households, whereas a capital income tax generates a negative wealth effect, pushing wealth-rich households to work. In this paper, we propose a new mechanism based on consumption quality, develop a model of this mechanism, and provide empirical validation for it. In our incomplete-markets model with quality choice, wealth-rich individuals keep working to purchase expensive, high-quality consumption.

Our paper is also related to the work studying how hours worked vary over time and across countries. Boppart and Krusell (2020) propose a utility function that admits falling hours along a balanced growth path, consistently with the historical, time-series evidence for several OECD countries.[4] Similarly, Restuccia and Vandenbroucke (2013) leverage the income effects on labor supply from Stone-Geary preferences to explain the decline in U.S. hours worked over the last century. Bick, Fuchs-Schündeln and Lagakos (2018) argue that income effects are responsible for the negative relationship between hours worked and GDP per capita. What distinguishes our work is the focus on the *cross-sectional* distribution of hours over wealth. Given the total hours worked in the economy, here we study how those hours are distributed across households with different wealth. Notably, we propose a theory that can reconcile the highly concentrated distribution of wealth with the observation that hours worked are more evenly distributed than wealth.

Finally, Jaimovich, Rebelo and Wong (2019) study the business cycle implications of the quality choice in a representative-agent model and a model in which heterogeneous agents achieve full insurance within a family that pools income and shares wealth across its members. Naturally, given their model setup, they do not examine the cross-sectional distribution of hours over wealth. However, they show that high-quality consumption is labor-intensive, which magnifies the impact of cyclical shocks on hours and output. Here, we ask a different question and learn a new mechanism. With uninsurable wage risk, a quality choice with non-homothetic preferences also has important implications for the cross-sectional distribution of hours. Further, quality choice acts as a self-insurance mechanism against adverse productivity shocks.

3. Quality Choice and Labor Supply

This section provides insight into how the quality choice affects labor supply, taking wealth as exogenously given. In the next Section 4, we embed the quality choice into an infinite-horizon model with uninsurable wage shocks, in which precautionary saving fuels wealth accumulation, giving rise to an *endogenous* wealth distribution.

3.1. A Standard Labor Supply Model

An individual has preferences $u(c) - Bh$, where $u(\cdot)$ is strictly increasing, concave, and twice continuously differentiable, c is consumption, B is the disutility of work, and $h \in$

[4] See also Cociuba, Prescott and Ueberfeldt (2018) for the U.S., and Bick, Fuchs-Schündeln and Lagakos (2018) and Bridgman, Duernecker and Herrendorf (2018) for cross-country evidence on hours worked.

$\{0,1\}$ is indivisible labor supply as in Rogerson (1988). The individual's problem is to maximize $u(c) - Bh$ by choosing to work ($h = 1$) or not to work ($h = 0$), subject to the budget constraint $c = wh + a$, where w is the wage and a is wealth. The value of working is $V^E = u(w + a) - B$, whereas the value of not-working is $V^N = u(a)$. The individual's decision is to work if $V^E > V^N$. The labor supply choice follows a reservation wage rule: for a given level of wealth a, there is a unique cutoff on the wage, w_R, such that if $w \leq w_R$, the individual does not work, otherwise if $w > w_R$, the individual works. Such a cutoff is implicitly determined by the indifference condition between working and not-working, $u(w_R + a) - B = u(a)$. Total differentiation of this condition gives that the reservation wage is monotonically increasing in wealth, $dw_R/da = [u'(a) - u'(w_R + a)] / u'(w_R + a) \geq 0$, since $u'(a) \geq u'(w_R + a)$ from the concavity of the utility function. Everything else equal, the larger the wealth, the higher the reservation wage, and the weaker the incentives to work.[5]

As shown in Figure 1, employment rates are nearly flat across wealth deciles. There are two aspects of this observation worth stressing. First, in the data, household heads with virtually zero wealth work "too little" compared to what the model predicts. This discrepancy comes from the assumption that, in the model, individuals out of work have no income to fund consumption, which is unrealistic. Allowing for government transfers targeted at the wealth-poor has been the typical approach in the literature to overcome this counterfactual model prediction. Second, the standard labor supply model predicts that work incentives taper off as wealth increases. What is needed is, then, a mechanism that flattens the cross-sectional relation between reservation wages and wealth. Next, we show that a quality choice in consumption can account under some conditions for the cross-sectional relationship between hours worked and wealth in the data.

3.2. A Labor Supply Model with Quality Choice

We modify the individual's labor supply problem by allowing for a quality choice in consumption. Preferences are $u(c, q) - Bh$ where q denotes the quality of consumption. The individual's problem is to maximize $u(c, q) - Bh$, subject to the budget constraint $p(q)c = wh + a$. To capture the idea that higher-quality versions are more expensive, we assume that $p'(q) \equiv dp(q)/dq \geq 0$.

[5]In Online Appendix B, we show that the same prediction holds in (i) a model with capital income taxes, (ii) a neoclassical growth model with wealth heterogeneity, and (iii) a version of the neoclassical growth model with idiosyncratic wage shocks, wealth heterogeneity, and complete markets.

Combining the first-order conditions (FOCs) for consumption quantity and quality gives the intratemporal condition

$$\frac{u_2(c,q)}{u_1(c,q)} = \frac{p'(q)c}{p(q)}, \tag{1}$$

where u_1 denotes the derivative of the utility function with respect to consumption, i.e., the marginal utility of consumption. Similarly, u_2 is the marginal utility of quality.

Total differentiation of (1) yields

$$\frac{dq}{dc} = \left[\frac{1}{c} - \left(\frac{u_{21}(c,q)}{u_2(c,q)} - \frac{u_{11}(c,q)}{u_1(c,q)}\right)\right] \Bigg/ \left[\frac{u_{22}(c,q)}{u_2(c,q)} - \frac{u_{12}(c,q)}{u_1(c,q)} - \left(\frac{p''(q)}{p'(q)} - \frac{p'(q)}{p(q)}\right)\right], \tag{2}$$

so that $dq/dc = 0$ if and only if

$$\frac{1}{c} - \frac{u_{21}(c,q)}{u_2(c,q)} + \frac{u_{11}(c,q)}{u_1(c,q)} = 0. \tag{3}$$

Condition (3) defines the class of admissible utility functions consistent with quality being invariant to the quantity of consumption.

Proposition 1 (Separable preferences) *Assume that preferences are separable in the quantity and quality of consumption, such that $u(c,q) = f(c) + g(q)$, where f and g are strictly increasing and concave, and twice continuously differentiable. Quality choice is invariant to the quantity of consumption, i.e., $dq/dc = 0$, if and only if the utility function is logarithmic in consumption:*

$$f(c) = \alpha \log(c),$$

where $\alpha > 0$ is an arbitrary constant.

Proof. See Online Appendix B. ■

Proposition 2 (Non-separable preferences) *Assume that preferences are non-separable in the quantity and quality of consumption, such that $u(c,q) = f(c)g(q)$, where f and g are strictly increasing and concave, and twice continuously differentiable. Quality choice is invariant to the quantity of consumption, i.e., $dq/dc = 0$, if and only if the marginal rate of substitution is proportional to consumption:*

$$MRS \equiv \frac{u_2(c,q)}{u_1(c,q)} \propto c.$$

Proof. See Online Appendix B. ■

To summarize, the requirement of quality to be a normal good so that higher-income individuals choose higher-quality goods imposes restrictions on preferences. That is, preferences are to be *non-homothetic*.

The reservation wage is implicitly determined by the indifference condition

$$u\left(\frac{w_R + a}{p(q_e)}, q_e\right) - B = u\left(\frac{a}{p(q_u)}, q_u\right), \tag{4}$$

where q_e and q_u indicate the quality choice if working, and not-working, respectively. Total differentiation of (4) gives

$$\left[\frac{u_1\left(c_e, q_e\right)}{p(q_e)}\right] dw_R = \left[\frac{u_1\left(c_u, q_u\right)}{p(q_u)} - \frac{u_1\left(c_e, q_e\right)}{p(q_e)}\right] da + \left[u_2\left(c_u, q_u\right) - \frac{p'(q_u)c_u}{p(q_u)} u_1\left(c_u, q_u\right)\right] dq$$
$$- \left[u_2\left(c_e, q_e\right) - \frac{p'(q_e)c_e}{p(q_e)} u_1\left(c_e, q_e\right)\right] dq, \tag{5}$$

where $c_e = (w_R + a)/p(q_e)$ and $c_u = a/p(q_u)$ indicate consumption if working, and not working, respectively. The intratemporal condition (1) implies that the last two terms on the right-hand side of (5) are equal to zero, such that

$$\frac{dw_R}{da} = \frac{p(q_e)}{p(q_u)} \cdot \frac{u_1\left(c_u, q_u\right)}{u_1\left(c_e, q_e\right)} - 1 \lesseqgtr 0. \tag{6}$$

The key insight from (6) is that the sign of the comparative statics depends on two channels. The first pertains to the available menu of quality-price bundles, as captured by the relative price term $p(q_e)/p(q_u)$. The second measures the extent to which the marginal utility of consumption changes based on the decision to work or not, as captured by the ratio of marginal utilities term $u_1\left(c_u, q_u\right)/u_1\left(c_e, q_e\right)$. In general, then, the reservation wage can be increasing or decreasing in wealth, or even invariant to wealth if the knife-edge condition $p(q_e)u_1\left(c_u, q_u\right) = p(q_u)u_1\left(c_e, q_e\right)$ holds.

There are, however, two cases in which we can provide a definite answer.

Proposition 3 (Irrelevance of quality choice #1) *If the utility function $u(c, q)$ defined over the quantity c and quality q of consumption is strictly increasing and concave, twice continuously differentiable, and it satisfies the restriction that*

$$\frac{1}{c} - \frac{u_{21}(c, q)}{u_2(c, q)} + \frac{u_{11}(c, q)}{u_1(c, q)} = 0,$$

then the reservation wage is monotonically increasing in wealth.

Proof. If $q_e = q_u = \bar{q}$, then $p(q_e) = p(q_u) = p(\bar{q})$, implying that $p(q_e)/p(q_u) = 1$. Since $u_1(a/p(\bar{q}), \bar{q}) \geq u_1((w_R + a)/p(\bar{q}), \bar{q})$ from the concavity of the utility function, equation (6) implies that $dw_R/da \geq 0$ for all $a \geq 0$. ■

Proposition 3 provides an important benchmark. In the case of homothetic preferences in which quality choice does not depend on the employment status, the reservation wage remains increasing in wealth, as in the standard labor supply model without the quality choice. In this sense, a quality choice is a necessary, not sufficient, condition to offset the negative wealth effect on labor supply.

Proposition 4 (Irrelevance of quality choice #2) *If the utility function $u(c, q)$ is separable in the quantity c and quality q of consumption (i.e., the marginal utility of consumption is invariant to quality) and prices are (weakly) increasing in quality, then the reservation wage is monotonically increasing in wealth.*

Proof. With a separable utility function, equation (6) becomes

$$\frac{dw_R}{da} = \frac{p(q_e)}{p(q_u)} \cdot \frac{u_1(c_u)}{u_1(c_e)} - 1 \geq 0. \tag{7}$$

If $p'(q) \geq 0$, then $p(q_e)/p(q_u) \geq 1$. With $c_e \geq c_u$, $u_1(c_u) \geq u_1(c_e)$ from the concavity of the utility function, such that $dw_R/da \geq 0$ for all $a \geq 0$. ■

Proposition 4 provides another important benchmark result. Insofar as consumption is a normal good, a form of non-separability between quality and quantity is needed to overturn the result of the reservation wage rising with wealth. That is, a higher quality must imply not only a higher utility but also a higher marginal utility of consumption, i.e., $u_{12}(c, q) > 0$. Note, however, that non-separability is only a necessary condition; the positive effect of quality on the marginal utility of consumption has to be strong enough to offset the relative price effect.

3.3. A Static General-Equilibrium Model with Quality Choice

A key theoretical insight from the previous subsection is that a quality choice with non-homothetic preferences is necessary, yet not sufficient condition to reproduce the nearly flat relationship between employment rates and wealth in the data. Crucially, functional form assumptions for utility and production functions, as well as the general equilibrium that determines the menu of quality-price bundles, are critical factors in determining the sign of the comparative statics of the reservation wage to wealth. In this subsection, we study a static general-equilibrium model featuring a quality choice with non-homothetic preferences and Cobb-Douglas production functions to clarify some of these issues.

On the consumption side, we amend the utility function in Jaimovich, Rebelo and Wong (2019) to allow for indivisible labor, and assume that preferences are described by $U = \left[q^{1-\theta}/\left(1-\theta\right)\right]\log(c) - Bh$, with $0 < \theta < 1$. The FOCs for consumption quantity and quality give the intratemporal condition:

$$\frac{p'(q)q}{p(q)} = (1-\theta)\log(c) = (1-\theta)\log\left(\frac{wh+a}{p(q)}\right). \tag{8}$$

On the production side, there are sectors producing consumption goods that differ by quality. Within each sector, perfectly competitive firms produce Y_q units of the final good of quality q using a Cobb-Douglas production function, $Y_q = K_q^{\alpha}\left(N_q/q\right)^{1-\alpha}$, with $0 < \alpha < 1$, where K_q and N_q are capital and labor, respectively. Dividing the FOCs for capital and labor gives the capital-labor ratio as independent of quality, $K_q/N_q = \left[\alpha/\left(1-\alpha\right)\right]w/R$, where w and R are the wage and capital rental rate, respectively. Next, using the expression for the capital-labor ratio and the FOC for labor, after rearranging terms, we obtain

$$p(q) = q^{1-\alpha}\left(\frac{w}{1-\alpha}\right)^{1-\alpha}\left(\frac{R}{\alpha}\right)^{\alpha} = Gq^{1-\alpha}, \tag{9}$$

where $G \equiv \left[w/(1-\alpha)\right]^{1-\alpha}\left(R/\alpha\right)^{\alpha}$. Note that (9) implies that prices are increasing in quality and that the price elasticity to quality is constant and equal to $1-\alpha$.

Using (8)-(9), we obtain that consumption quality, $q = (wh+a)^{\frac{1}{1-\alpha}}/(e^{\frac{1}{1-\theta}}G^{\frac{1}{1-\alpha}})$, and the unit prices, $p(q) = (wh+a)/e^{\frac{1-\alpha}{1-\theta}}$, are increasing in earnings and wealth, and that consumption quantity, $c = e^{\frac{1-\alpha}{1-\theta}}$, is constant. The model yields sharp predictions on Engel curves: the quantity Engel curve is a flat line, whereas the quality Engel curve is linear with a unitary price elasticity to income.

Finally, the individual's indifference condition between working and not-working gives the reservation wage,

$$w_R = \left[a^{\frac{1-\theta}{1-\alpha}} + \frac{(1-\theta)^2}{1-\alpha}G^{\frac{1-\theta}{1-\alpha}}eB\right]^{\frac{1-\alpha}{1-\theta}} - a. \tag{10}$$

First, if the elasticity of utility to quality equals the price elasticity to quality ($\theta = \alpha$), the reservation wage is independent of wealth. Note that while in the Cobb-Douglas case, this result depends on a knife-edge condition on parameters, this is not true with CES production functions. Generally, the price elasticity to quality is an endogenous object determined alongside equilibrium allocations and prices. This comparative statics result thus provides a helpful benchmark. Insofar as utility and (equilibrium) price elasticities to quality are roughly the same, the reservation wage is insensitive to changes in wealth, implying that employment rates do not fall steeply with wealth.

Second, if the elasticity of utility to quality is larger than the price elasticity to quality ($\theta < \alpha$), the reservation wage decreases in wealth. The higher the wealth, the lower the reservation wage, and the higher the likelihood of working. Such a negative relationship between reservation wages and wealth implies that employment rates are increasing in wealth. If instead, the elasticity of utility to quality is smaller than the price elasticity ($\theta > \alpha$), the reservation wage is increasing in wealth. The larger the wealth, the higher the reservation wage, implying that employment rates are decreasing in wealth. Again, something we do not see in the data.

While valuable for analytical insight, the Cobb-Douglas structure puts restrictions on the shape of quality Engel curves that do not hold in the data, such as a unitary price elasticity to income. Hence, in Section 4, we assume a CES production structure in which the price elasticity to quality becomes an equilibrium object determined alongside wages, rental rate, and the level of quality itself.

4. Incomplete-Markets Model with Quality Choice

In this section, we develop a quantitative theory of how households make consumption and labor supply decisions and how these decisions impact wealth accumulation and, thereby, the relationship between wealth and hours worked across households.

4.1. Environment

Time is discrete and continues forever. There is a continuum of measure one of infinitely-lived households, who choose consumption quantity and quality, whether to work, and how much to save in the face of idiosyncratic shocks. Households have one unit of time per period, which yields z_t units of labor services, where z_t is *i.i.d.* across households.

Preferences and budget constraint. Before describing preferences and budget sets, it is helpful to discuss the household's choice problem. At any point in time, households face a continuum of quality-price bundles $\{q_t, p_t(q_t)\}$ from which to choose, where unit prices $p_t(q_t)$ are functions of quality levels q_t. Households can choose only *one* bundle from those available, whereas they can consume any quantity $c_{q,t}$ of the consumption good of quality q_t. Henceforth, abusing notation slightly, we use p_t to denote the price function $p_t(q_t)$ and c_t to denote the quantity $c_{q,t}$.

Preferences over streams of consumption and hours worked are described by $U = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[u(c_t, q_t) - v(h_t)\right]$, where $\mathbb{E}_0$ is the mathematical expectation at $t = 0$, $0 < \beta < 1$ is the time discount factor, and the period utility function $u(\cdot)$ is strictly increasing, concave, and twice continuously differentiable in both arguments. Labor is indivisible, and $v(\cdot)$ satisfies $v\left(\bar{h}\right) = B\bar{h}$ when $h_t = \bar{h} > 0$, and $v(0) = 0$ when $h_t = 0$.

Household's expenditures are purchases of the consumption good, $(1 + \tau_c)p_t c_t$, where τ_c is a constant consumption tax rate, and a one-period risk-free asset, a_{t+1}, that earns the after-tax rate of return $(1 - \tau_k)r_t$, where τ_k is a constant capital tax rate. Income comes from three sources: (i) after-tax labor income or earnings, $z_t w_t h_t - \mathcal{T}(z_t w_t h_t)$, where w_t is the hourly wage per efficiency units of labor z_t, and $\mathcal{T}(z_t w_t h_t)$ is labor income taxes calculated based on the tax function $\mathcal{T}(\cdot)$, whose features we discuss further below; (ii) asset income, $(1 - \tau_k)r_t a_t$; and (iii) means-tested government transfers, $T(a_t, z_t w_t \bar{h}) \geq 0$, which depend on assets and labor earnings in a way we specify below.

The household's budget constraint is

$$a_{t+1} = a_t + (1 - \tau_k)r_t a_t + z_t w_t h_t - \mathcal{T}(z_t w_t h_t) - (1 + \tau_c)p_t c_t + \mathbb{1}_{(h_t=0)} T(a_t, z_t w_t \bar{h}), \quad (11)$$

where the indicator function $\mathbb{1}_{(h_t=0)}$ equals one if the household is nonemployed, and zero if employed. We also impose an exogenous limit on how much the household can borrow, i.e., $a_t \geq \underline{a}$, with $\underline{a} \leq 0$.

Technology. The production technology for the consumption good of quality q is CES with capital, $K_{q,t}$, and labor services, $N_{q,t}$, as inputs:

$$Y_{q,t} = \left[\alpha K_{q,t}^{\rho} + (1 - \alpha)\left(N_{q,t}/q\right)^{\rho}\right]^{1/\rho}, \quad 0 < \alpha < 1, \quad \rho \leq 1. \quad (12)$$

The parameter ρ governs the degree of substitutability between capital and labor. For $\rho = 0$, the production function reduces to a Cobb-Douglas, $Y_{q,t} = K_{q,t}^{\alpha}\left(N_{q,t}/q\right)^{1-\alpha}$, and the elasticity of substitution is one. For $\rho < 0$, there is less substitution between capital and labor than in the Cobb-Douglas case. For future reference, we note that when $\rho < 0$, labor intensity increases in quality, i.e., the production of higher-quality goods displays lower capital-labor ratios.

Similarly, the production technology for the investment good is

$$X_t = \left[\alpha K_{I,t}^{\rho} + (1-\alpha)(N_{I,t}/q_I)^{\rho}\right]^{1/\rho}, \tag{13}$$

where q_I is the normalized quality of the investment good, and $K_{I,t}$ and $N_{I,t}$ are capital and labor services in the investment sector, respectively.[6]

Government. The government collects taxes on consumption expenditures and asset income with proportional tax rates $0 \leq \tau_c \leq 1$ and $0 \leq \tau_k \leq 1$, respectively, and labor income via the tax function $\mathcal{T}(y_t; \varphi_0, \varphi_1)$, where y_t is pre-tax labor income and (φ_0, φ_1) are two parameters governing respectively the level and progressivity of taxation in a way we explain further in Section 5. Outlays are government consumption, G_t, and transfers, T_t. We assume that the government balances the budget by adjusting φ_0 so that tax revenues equal government outlays.

Transfers $T_t \geq 0$ are means-tested by the household's wealth and after-tax earnings and guarantee a minimum $\bar{c}$ of consumption expenditures to non-working households. Absent these transfers, households with zero wealth would necessarily work to finance consumption regardless of how low their productivity is. According to this program, non-working households receive $\bar{c}$ net of what they could afford by selling off their wealth, and their potential after-tax labor earnings:

$$T(a_t, z_t w_t \bar{h}) = \max\{0, \bar{c} - [a_t + (1-\tau_k) r_t a_t]\, \mathbb{1}_{(a_t>0)} - \left[z_t w_t \bar{h} - \mathcal{T}(z_t w_t \bar{h})\right]\}. \tag{14}$$

4.2. The Household's Problem

We formulate the household problem in recursive form and use primes to denote next period variables. All information necessary for optimal decision making at a particular point in time is summarized by the state vector (a, z), where a and z are the individual state variables. We omit the aggregate state variables from the state vector as we focus on the stationary equilibrium of the model, in which aggregate variables are constant.

[6] In computing the model's equilibrium, we set q_I equal to the average quality of the consumption good.

The value function of a household that begins the period with the state vector (a, z) is $V(a,z) = \max\left\{V^E(a,z), V^N(a,z)\right\}$, where $V^E(a,z)$ and $V^N(a,z)$ are the value functions conditional on working and not working, respectively. If the household chooses to work because $V^E(a,z) > V^N(a,z)$, it solves the dynamic problem:

$$V^E(a,z) = \max_{c,\,q,\,a'} \left\{u(c,q) - B\bar{h} + \beta\mathbb{E}\left[V(a',z')|z\right]\right\}; \tag{15}$$

$$\text{s.t.} \quad a' = a + (1-\tau_k)ra + zw\bar{h} - \mathcal{T}(zw\bar{h}) - (1+\tau_c)p(q)c; \tag{16}$$

$$a' \geq \underline{a}. \tag{17}$$

If the household instead chooses not to work because $V^N(a,z) \geq V^E(a,z)$, it solves the dynamic problem:

$$V^N(a,z) = \max_{c,\,q,\,a'} \left\{u(c,q) + \beta\mathbb{E}\left[V(a',z')|z\right]\right\}; \tag{18}$$

$$\text{s.t.} \quad a' = a + (1-\tau_k)ra - (1+\tau_c)p(q)c + T(a, zw\bar{h}); \tag{19}$$

$$a' \geq \underline{a}. \tag{20}$$

As in the static model in Section 3, the household's decision problem involves a trade-off between consumption quantity and quality. Combining the FOCs with respect to consumption quantity and quality gives the intratemporal condition that captures this quality-quantity trade-off:

$$\frac{u_2(c,q)}{u_1(c,q)} = \frac{p'(q)c}{p(q)}, \tag{21}$$

where $p'(q) \equiv dp(q)/dq$.

4.3. Production

Production of consumption and investment goods takes place in perfectly competitive markets. We assume capital and labor can freely move across sectors, so wages and capital rental rates are equalized. We begin by describing the firm problem in the consumption sector and then turn to the investment sector.

Consumption sector. Firms in the consumption sector maximize profits, $\Pi_q \equiv p(q)Y_q - RK_q - wN_q$, taking the wage, w, capital rental rate, R, and output price, $p(q)$, as given, subject to the technology (12). Combining the FOCs with respect to capital and labor gives the capital-labor ratio in the consumption sector:

$$\frac{K_q}{N_q} = \left[\left(\frac{\alpha}{1-\alpha}\right)\frac{w}{R}\right]^{\frac{1}{1-\rho}} q^{\frac{\rho}{1-\rho}}. \tag{22}$$

In the Cobb-Douglas case ($\rho = 0$), the capital-labor ratio in (22) is independent of quality. If $0 < \rho < 1$, the capital-labor ratio increases in quality – quality is capital intensive. If instead $\rho < 0$, the capital-labor ratio decreases in quality – quality is labor intensive.

Using the capital-labor ratio (22), and rearranging terms, we obtain

$$p(q) = \left[(1-\alpha)^{1/(1-\rho)} (wq)^{\frac{\rho}{\rho-1}} + \alpha^{1/(1-\rho)} R^{\frac{\rho}{\rho-1}}\right]^{\frac{\rho-1}{\rho}}. \tag{23}$$

The price function (23) makes evident that the menu of quality-price bundles available for purchase is an equilibrium object, as it depends on the market-clearing values of w and R. From (23), unit prices are increasing in quality:

$$p'(q) = \frac{(1-\alpha)^{1/(1-\rho)} w^{\frac{\rho}{\rho-1}} q^{\frac{1}{\rho-1}}}{\left[(1-\alpha)^{1/(1-\rho)} (wq)^{\frac{\rho}{\rho-1}} + \alpha^{1/(1-\rho)} R^{\frac{\rho}{\rho-1}}\right]^{\frac{1}{\rho}}} \geq 0. \tag{24}$$

Note also that the price elasticity to quality $p'(q)q/p(q)$ is an equilibrium object, too, and depends on w, R, and the level of quality itself. For $0 < \rho < 1$, the price elasticity is decreasing in quality, whereas, for $\rho < 0$, it is increasing in quality. In the Cobb-Douglas case, the price elasticity to quality is constant and equal to $1 - \alpha$.

Investment sector. We choose the investment good to be the numéraire, so its price is normalized to one. Firms in the investment sector maximize profits, $\Pi_I \equiv X - RK_I - wN_I$, subject to the production technology (13). Combining the FOCs with respect to capital and labor gives the capital-labor ratio in the investment sector:

$$\frac{K_I}{N_I} = \left[\left(\frac{\alpha}{1-\alpha}\right)\frac{w}{R}\right]^{\frac{1}{1-\rho}} q_I^{\frac{\rho}{1-\rho}}. \tag{25}$$

4.4. Competitive Equilibrium

We consider a stationary equilibrium in which aggregate variables are constant.

Equilibrium definition. A Recursive Competitive Equilibrium (RCE) consists of a set of value functions, $\{V(a,z), V^E(a,z), V^N(a,z)\}$, a set of decision rules for the quantity and quality of consumption, asset holdings, and labor supply, $\{c_q(a,z), q(a,z), a'(a,z), h(a,z)\}$, aggregate inputs in the consumption sector, $\{K_C, N_C\}$, inputs in the investment sector, $\{K_I, N_I\}$, factor prices, $\{w, R\}$, unit prices of different qualities, $\{p(q(a,z))\}$, government policy, $\{G, T(a,z), \tau_c, \tau_k, \mathcal{T}(zwh(a,z))\}$, and a stationary distribution $\lambda(a,z)$ induced by the stochastic process for z and the decision rule for asset holdings $a'(a,z)$, such that:

1. Individual decision rules solve Bellman equations.
2. Firms maximize profits.
3. The asset market clears: $\int a d\lambda = K$, where $K = K_C + K_I$ is aggregate capital, and K_C and K_I are the capital stocks in the consumption and investment sector, respectively.
4. The labor market clears: $\int zh(a,z)d\lambda = N$, where $N = N_C + N_I$ is the aggregate labor input, and N_C and N_I are labor inputs in the consumption and investment sectors, respectively. Note that N is aggregate efficiency-weighted hours. Aggregate hours worked are $H = \int h(a,z)d\lambda$.
5. The government budget constraint is balanced:
$$G + \int T(a, zw\bar{h})d\lambda = \tau_c \int p(q(a,z))c_q(a,z)d\lambda + \tau_k r \int a d\lambda + \int \mathcal{T}(zwh(a,z))d\lambda.$$
6. The market for each quality level q clears: $\int c_q(a,z)d\lambda = Y_q$.
7. The goods market clears:
$$\begin{aligned}&\int \left[p(q(a,z))c_q(a,z) + a'(a,z)\right] d\lambda \\ &= \int p(q(a,z)) \left[F_1\left(K_q, \frac{N_q}{q}\right) a + F_2\left(K_q, \frac{N_q}{q}\right) zh(a,z)\right] d\lambda + (1-\delta)K,\end{aligned}$$
where $F_1\left(K_q, N_q/q\right) \equiv \partial Y_q/\partial K_q$ and $F_2\left(K_q, N_q/q\right) \equiv \partial Y_q/\partial N_q$.
8. The stock of capital evolves according to $K' = (1-\delta)K + X$, where $X = \delta K$ and $0 < \delta < 1$ is the capital depreciation rate.

5. Parameterization

We now discuss the model's calibration describing preferences, technology, taxes, and transfers. In dynamic general equilibrium models, none of the parameters has a one-to-one relationship to a specific moment. Nonetheless, describing the calibration procedure as a few distinct steps is helpful. First, we exogenously set some parameters directly from the data or based on values in the literature. Second, we jointly calibrate the remaining ones to match a select number of data moments.

5.1. Preferences

Time discount factor. Decisions in the model take place at a quarterly frequency. As customary in the literature, we calibrate the time discount factor to replicate the annual after-tax rate of return on capital of 4% (Gomme, Ravikumar and Rupert, 2011; McGrattan and Prescott, 2003). This procedure yields a value for β of 0.975.

Disutility of work. We normalize labor supply to $\bar{h} = 1$ and calibrate the disutility of work parameter B to match the aggregate employment rate of 80%. This procedure yields a value for B of 1.78. Alternatively, one could set $\bar{h} = 1/3$ so that a working household spends one-third of available time at work, as in the data, and then multiply the value of the disutility of work B by three.

Curvature in quality. In Section 3, we derived theoretical results ruling out homothetic utility function specifications. Further, we established that a non-separability in quality is necessary to restore work incentives for wealthy individuals. Given these requirements, we adopt the utility function in Jaimovich, Rebelo and Wong (2019), amended to allow for a labor supply choice along the extensive margin, $\left[q^{1-\theta}/\left(1-\theta\right)\right]\log(c) - Bh$. This functional form specification has several appealing properties. First, the marginal utility of consumption is increasing in quality. This property is critical for the model to generate work incentives that are strong enough to offset the negative wealth effect on labor supply. To see this, consider the basic insight from standard labor supply theory again. In deciding whether to work or not, an individual trades off the utility gain from working due to the additional consumption one can afford with the disutility of work. As consumption increases with wealth and the marginal utility decreases, the utility gain from working for a wealth-rich individual is necessarily smaller than for a wealth-poor individual. As a result, the willingness to supply labor decreases with wealth. Second, note that if we set $q = 1$, the preferences above nest the utility function $\log(c) - Bh$, which is one of macroeconomics's most widely used specifications.

Individuals' preferences over quality are described by θ, which controls the curvature of the utility function in quality, and, more specifically, how fast the marginal utility of consumption (MUC) increases with quality, as captured by $u_{12}(c, q) = q^{-\theta}/c \geq 0$. To pin down the value of θ, we leverage the structure of the model and calibrate θ jointly with other parameter values in the context of a simulated method of moments (SMM) exercise.

Note that θ directly impinges on quality choices via the intratemporal condition (21). For any level of assets, the model generates a positive relationship between income and quality, which implies a positive relationship between income and unit prices. Everything else equal, the equilibrium dispersion in unit prices critically depends on the value of θ. With this in mind, we require the model to reproduce the percent difference in average unit prices between the fourth and the first household income quartile in the data. In Section 7, we estimate the empirical relationship between unit prices and income using durables and housing expenditures from CEX and homescan data from Nielsen. Our estimates yield a target of 0.54, the weighted average of the unit price differences between the fourth and first income quartiles in the two datasets, calculated using expenditure shares in CEX as weights.[7] This procedure yields a value for θ of 0.73.

5.2. Technology

We set δ to 1.5% per quarter, which yields a capital depreciation rate of 6% a year. The capital income share α is 36% (see, e.g., Koh, Santaeulàlia-Llopis and Zheng, 2020). The parameter ρ determines the elasticity of substitution between capital and labor, $1/(1 - \rho)$. We set $\rho = -0.5$, as in Jaimovich, Rebelo and Wong (2019), such that the capital-labor elasticity is 0.67, implying that capital and labor are complements. This value for the elasticity is in line with the range of empirical estimates 0.5-0.7 in the literature (see, e.g., Oberfield and Raval, 2021).

[7] Using the crosswalk between CEX and Nielsen from Coibion, Gorodnichenko and Koustas (2021), we associate a weight of 0.53 to durables, 0.19 to housing, and 0.28 to the product categories in Nielsen. The targeted moment is then calculated as $0.53 \times 0.677 + 0.19 \times 0.607 + 0.28 \times 0.241 \approx 0.54$, where the values of 0.677, 0.607, and 0.241 are the percent differences in average unit prices between the fourth and the first income quartile for durables and housing in CEX and Nielsen, respectively. See Section 7 for more details on sample selection and estimation.

5.3. Earnings Heterogeneity and Borrowing Limit

Wage shocks. The calibration of the stochastic process for the idiosyncratic productivity shock follows a two step procedure. First, we assume an AR(1) process in logs:

$$\log(z_{t+1}) = \rho_z \log(z_t) + \sigma_z \epsilon_{t+1}, \tag{26}$$

where the parameters $0 \leq \rho_z \leq 1$ and $\sigma_z \geq 0$ govern the shocks' persistence and volatility, respectively. Based on Floden and Lindé (2001), we set the persistence of the productivity shock ρ_z to 0.975, and the standard deviation of the innovation to the productivity shock σ_z to 0.165. (See Online Appendix C for the grid and the matrix of transition probabilities of the discretized productivity process.)

Second, based on Castañeda, Díaz-Giménez and Ríos-Rull (2003), we allow for the realization of an extreme productivity outcome and that from that extreme productivity outcome, there is a nontrivial probability of a large fall in productivity. The combination of these features makes the highest earners have a significant demand for precautionary saving. Operationally, we introduce an additional productivity state, $z^{\max}$, that can be reached only from the second and third highest states with the same probability. This gives three additional parameters: $z^{\max}$, π^{up}, and π^{stay}, where π^{up} is the probability that z moves to $z^{\max}$, and π^{stay} is the probability that z remains at $z^{\max}$. We calibrate these three parameters for the model to match three data moments: (i) the wealth share of the top wealth decile (65.53%); (ii) the earnings share of the top earnings decile (35.87%); (iii) the earnings share of the top 1% of the earnings distribution (11.76%). This procedure gives $z^{\max} = 18.53$, $\pi^{\text{up}} = 0.42$, and $\pi^{\text{stay}} = 0.92$. This parametrization yields an equilibrium wealth dispersion of a magnitude comparable to that in U.S. data by heightening the precautionary saving motive of high earners.

Borrowing limit. Finally, we pin down the exogenous borrowing limit $\underline{a}$ so that the model reproduces the share of households with negative assets in the PSID of 13%. This procedure yields a value for $\underline{a}$ equivalent to -3% of total output.

5.4. The Tax-Transfer System

The tax and transfer system captures salient features of the U.S. government. We specify government consumption, G, as a share g of total output and set g to 21%, which is the average share of government purchases of goods and services in gross domestic product for the post-World War II period in the United States. We set the consumption tax rate τ_c to 8% and the capital tax rate τ_k to 35% based on our calculations for 1948-2020. These are average effective tax rates calculated using the method of Mendoza, Razin and Tesar (1994), which uses National Income and Product Account data and aggregates all levels of the government (federal, state, and local) into one government sector.

To capture the progressivity in the U.S. tax code, we adopt a parametric specification of the tax system according to which labor taxes are

$$\mathcal{T}(y) = y - \varphi_0 y^{1-\varphi_1}, \tag{27}$$

where y is pre-tax labor income, zwh, φ_0 governs the level of taxation, and φ_1 captures the extent of tax progressivity. When $\varphi_1 = 0$, the tax rate is constant and equal to $1 - \varphi_0$. When $\varphi_1 = 1$, the tax system implies full redistribution, such that after-tax labor income equals φ_0 for any level of pre-tax income. For $0 < \varphi_1 < 1$, the tax system is progressive, and the marginal tax rates are monotonically increasing in pre-tax income. We set φ_1 equal to 0.09 based on Heathcote, Storesletten and Violante (2020, Table 2), and vary φ_0 to balance the government budget. This procedure gives a value for φ_0 of 0.87.

The tax function in (27) has a long tradition in public economics; see, e.g., Musgrave (1959), Jakobsson (1976), and Kakwani (1977), and, more recently, Benabou (2002), Guner, Kaygusuz and Ventura (2014), and Heathcote, Storesletten and Violante (2017). Beyond its analytical tractability, it provides a good approximation of the complex U.S. tax and transfer system, except at the bottom of the income distribution, where marginal tax rates can be quite high due to the phasing out of means-tested programs. Also, note that the tax function has no floor for disposable income (the after-tax income of households with zero pre-tax income is zero). In the United States, however, several programs guarantee such a disposable income floor: Medicare, Medicaid, Supplemental Nutrition Assistance Program (SNAP), Temporary Assistance for Needy Families (TANF), and Unemployment Insurance (UI). For these reasons, the tax function's log-linear fit worsens in the income distribution's bottom decile.

Finally, to capture the high marginal tax rates at the bottom of the income distribution, we implement a simple means-tested transfer program whose generosity is parametrized by $\bar{c}$ as specified in (14). We set the value of $\bar{c}$ for the model to reproduce the average transfers-to-income ratio of 16.3% for the lowest wealth quintile in PSID. This approach yields a value for $\bar{c}$ of 74.95, equal to 39% of the wage rate, w.[8]

[8] In the model without the quality choice, $\bar{c} = 0.66$, equal to 31% of w.

6. Properties of the Calibrated Economy

We now discuss the implications for equilibrium prices and allocations. To gauge the role of the quality choice, we contrast the model's predictions for employment rates with those of the standard model without the quality choice. To compare the two versions of the model with and without quality on an equal footing, we have calibrated both economies to the same data moments.[9]

6.1. The Distributions of Employment, Earnings, and Wealth

In the United States, the distributions of earnings and, especially, that of household wealth are very concentrated and skewed to the right (see, e.g., Castañeda, Díaz-Giménez and Ríos-Rull, 2003; Díaz-Giménez, Glover and Ríos-Rull, 2011; Kuhn and Ríos-Rull, 2016). Figure 2 shows that the model does a good job of accounting for these phenomena.

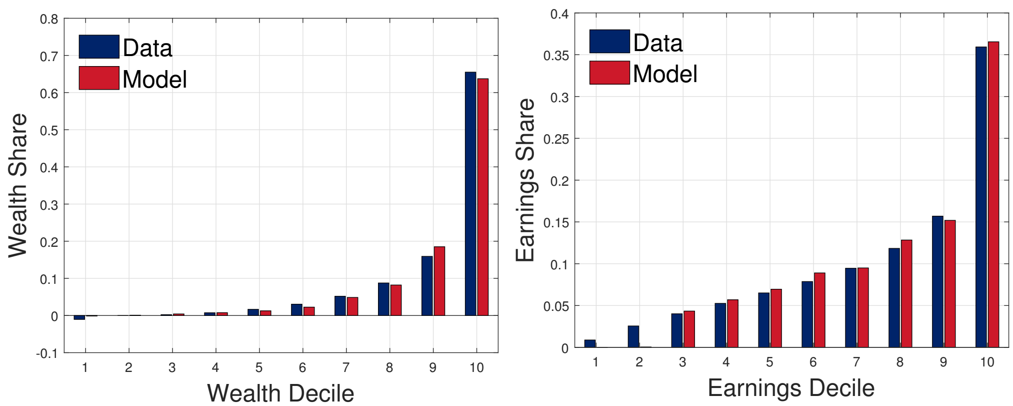

Figure 2: Earnings and Wealth Distributions – Model vs. Data

Notes: The figure shows the distributions of earnings and wealth in the model and in the data. Data are from the biannual 2001-2015 PSID waves for household heads of 25-65 years old. Wealth is total assets minus total liabilities at the household level. See Online Appendix A for details on variables' definitions and construction.

Figure 3 shows that the equilibrium cross-sectional distribution of employment rates by wealth implied by the model lines up very well with its empirical counterpart in the data. Notably, the model can account for the fact that the employment rates of wealth-rich households are nearly as high as those of wealth-poor households. In contrast, in the

[9]See Online Appendix C for details on the solution method used to compute the equilibrium of the model and Table C.1 for summary statistics on how the models with and without quality perform regarding targeted moments.

version of the model without quality, the employment rates of wealth-rich households are considerably lower than those of wealth-poor households, which again reflects the fact that the wealth effect on labor supply is too strong in the standard model. In the model with the quality choice and non-homothetic preferences, a significant fraction of wealthy households keeps on working to purchase expensive, high-quality versions of the consumption good.

The bottom of the wealth distribution. At the lower end of the wealth distribution, in the model, as in the data, households have zero or negative wealth. It is helpful to discuss the case of zero and negative wealth separately. In the case of zero or near-zero wealth, the standard labor supply model implies an employment rate of one. The intuition for this prediction is straightforward. If a household has no wealth, and so no asset income, and there are no government transfers, the only available option is to work at the ongoing wage to finance consumption. Note that this mechanism is at play in models with and without the quality choice. In the data, however, the employment rates for households with zero wealth are less than one; further, they are comparable to or even lower than those of households with nontrivial wealth.

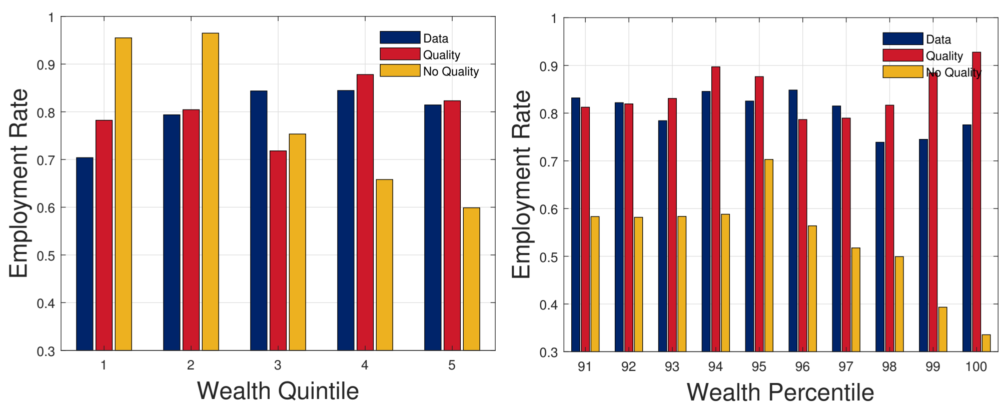

Figure 3: Employment Rates by Wealth – Model vs. Data

Notes: The figure shows the distribution of employment rates by wealth quintiles (left panel) and for the top 10% of the wealth distribution (right panel) in the models with and without the quality choice and in the data. Data are from the biannual 2001-2015 PSID waves for household heads of 25-65 years old. Wealth is total assets minus total liabilities at the household level. See Online Appendix A for details on variables' definitions and construction.

To tackle this issue, a standard approach in the literature is to allow for government transfers that guarantee a minimum level of consumption. These transfers go naturally to the wealth-poor in the model, thus capturing the number of transfer programs that direct resources to needy U.S. households (e.g., Medicare, Medicaid, SNAP, TANF, and UI). In equilibrium, such means-tested transfers significantly mitigate the precautionary motive of labor supply, implying that some households do not work despite having virtually any wealth, consistent with the data. Given the importance of transfers at the very bottom of the wealth distribution, we required the calibrated model to reproduce the 16.3% transfer-to-income ratio for the first wealth quintile in the PSID. In this respect, the model's ability to generate employment rates for wealth-poor households that are far below one, as in the data, is not an artifact of implausibly large transfers but rather a successful test of the theory. In contrast, the model without the quality choice generates unrealistic employment rates that are very close to one, in spite of being calibrated to match the same target for the transfer-to-income ratio.

In the case of negative wealth, it is important to distinguish between households that are against the borrowing constraint from those that retain the ability to borrow. In the former case, households are liquidity-constrained and adjust labor supply to smooth consumption. This self-insurance mechanism works in both models with and without the quality choice. In the model with quality choice, however, the constrained, wealth-poor households can cut back on quality to prevent or mitigate the drop in consumption. In this sense, quality choice provides additional self-insurance against adverse wage shocks, a mechanism absent in the standard model.

To be sure, labor supply and quality choices interact with borrowing decisions. For example, in response to a negative wage shock, a household can borrow in anticipation of future positive wage shocks, thus allowing for consumption smoothing above and beyond that coming from labor supply and quality choices. Indeed, to the extent that wealth-poor households are permitted to borrow and wage shocks are highly temporary, labor supply and quality choices are at least partly insulated from productivity shocks. By contrast, if wage shocks are very persistent, or in the extreme case, permanent, a household with negative wealth must work to repay debt.

The top of the wealth distribution. The right panel of Figure 3 zooms into the top 10% percent of the wealth distribution. In the data, employment rates remain remarkably stable at approximately 80%. As the figure shows, the differences between the two models with and without the quality choice are striking. Notably, the model with quality reproduces the wealth-hours profile in the data remarkably well. In contrast, in the model without quality, employment rates take a nosedive, which is at odds with the data. In the model without quality, the employment rate is as low as 35% at the top 1% percent of the wealth distribution, which contrasts with the roughly 75% employment rate in the data. Work incentives remain high even at the top of the wealth distribution in the model with quality, which contrasts sharply with the model without quality.

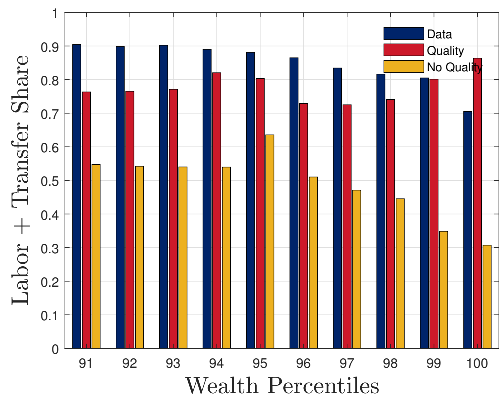

Figure 4: Labor and Transfer Income Shares by Wealth – Model vs. Data

Notes: The figure shows the shares of labor income and transfers in total income for the top 10% of the wealth distribution in the models with and without the quality choice and in the data. Data are from the biannual 2001-2015 PSID waves for household heads of 25-65 years old. Wealth is total assets minus total liabilities at the household level. See Online Appendix A for details on variables' definitions and construction.

As shown by Figure 4, the model with quality also accounts for the observation that labor earnings remain an important share of total income at the top of the wealth distribution, which is consistent with the evidence from the Survey of Consumer Finances and tax data for the United States (Kuhn, Schularick and Steins, 2020; Smith et al., 2019).

6.2. Consumption versus Expenditures

Here we discuss how consumption quantity, quality, and unit prices vary by wealth.

Consumption expenditures in the model. Figure 5 shows the equilibrium relationship between consumption quality and unit prices in the model, which traces the menu of quality-price bundles available to households. Not surprisingly, prices are increasing in quality: households are willing to pay more for higher-quality consumption. As evident from the figure, prices are approximately linear in quality. We stress that the slope of this quality-price relation is an equilibrium object, which depends on the technological parameters (α, ρ) and the equilibrium wage and capital rental rate. Hence, any change in the environment, for example, an unexpected tax policy change that causes a new equilibrium level of w and R, implies a shift in the price function and, thereby, a change in relative prices.

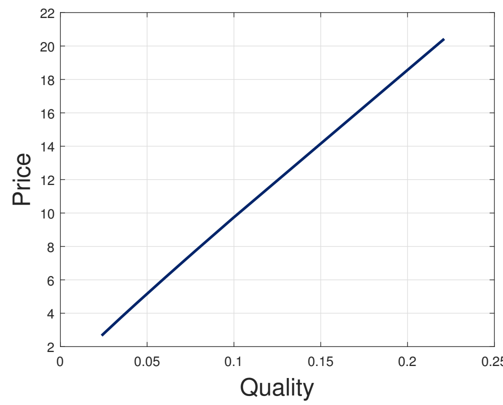

Figure 5: Equilibrium Price Function

Notes: The figure shows the relationship between prices and quality in the calibrated model.

Figure 6 shows consumption quantity and consumption expenditure, and unit prices by wealth in the calibrated model with the quality choice. A striking result emerges. While consumption quantity increases with wealth, it increases by less than expenditures. Most of the inequality in consumption expenditures in the model comes from differences in consumption quality and the higher unit prices paid by wealthy households.[10]

Figure 7 shows total consumption quantity and consumption expenditures shares by wealth, as implied by the calibrated models with and without the quality choice.[11] (In the model without quality, consumption quantity and expenditures are by construction identical.) In the model with quality, consumption quantity shares are nearly flat across the wealth distribution. By contrast, consumption expenditure shares are highly concentrated, with the top decile accounting for approximately 30% of total expenditures. Again, such inequality in consumption expenditures comes from wealthier households consuming roughly the same amount of consumption but purchasing more expensive, higher-quality versions.

[10]In the model with Cobb-Douglas production functions, consumption quantity is constant across wealth so that differences in consumption expenditures come entirely from differences in unit prices.

[11]To gauge consumption expenditures inequality relative to the data, the 90-10 percentile ratios in the models with and without the quality choice and PSID data are 6.02, 3.78, and 5.46, respectively. The 90-50 percentile ratios are 1.86, 1.71, and 2.10.

In the standard model without quality, consumption quantity rises considerably with wealth. Notably, the first decile of the wealth distribution accounts for 5% of total consumption, whereas the top decile accounts for nearly 26% of the total. This property makes the standard model at odds with the cross-sectional empirical evidence on employment rates and wealth. A pattern of increasing consumption quantities implies decreasing marginal utilities of consumption, which in turn leads to employment rates falling sharply with wealth.

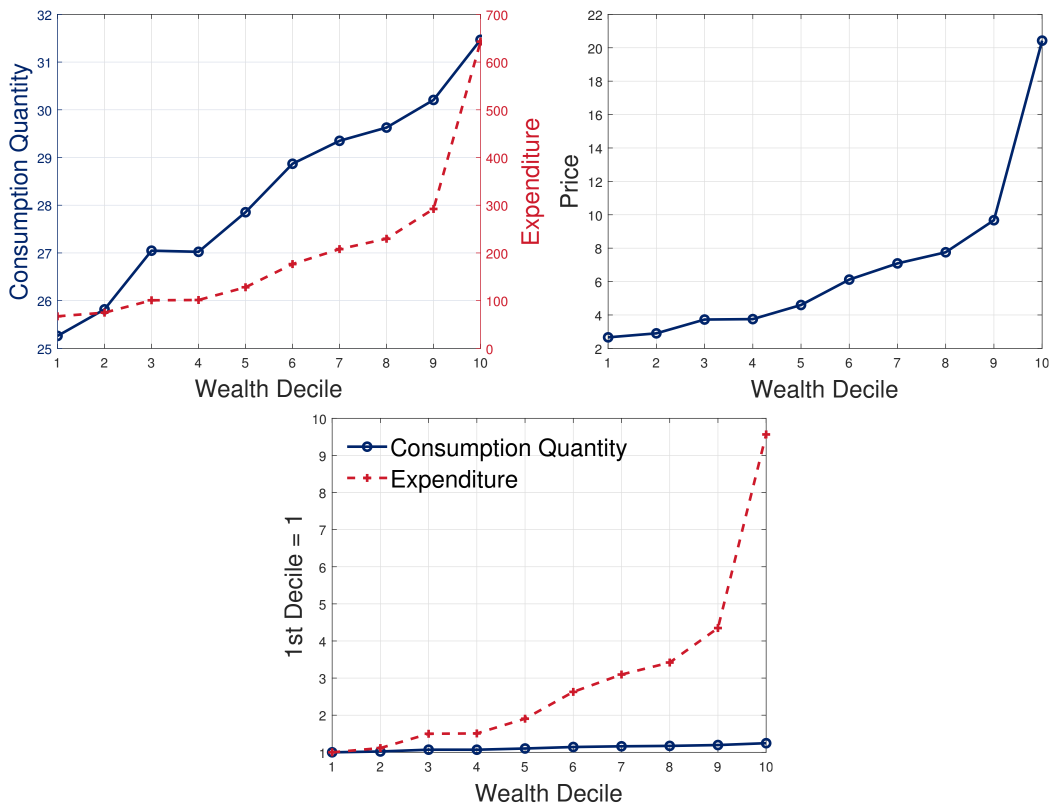

Figure 6: Consumption and Unit Prices by Wealth

Notes: The figure shows consumption quantity and consumption expenditures (top left panel), and unit prices (top right panel) by wealth deciles in the calibrated model with the quality choice. To facilitate comparison, the bottom panel shows consumption and expenditures by wealth, where we normalize the values of both variables in the first wealth decile to one.

Two important insights emerge from these results. First, the calibrated versions of the two models with and without quality have different implications for how consumption quantity varies by wealth. In the standard model, consumption rises with wealth, whereas in the model with quality, it is virtually flat in wealth. Such a difference in consumption allocations is why the model with quality accounts for the nearly flat distribution of employment rates across wealth deciles, whereas the standard model without quality cannot.

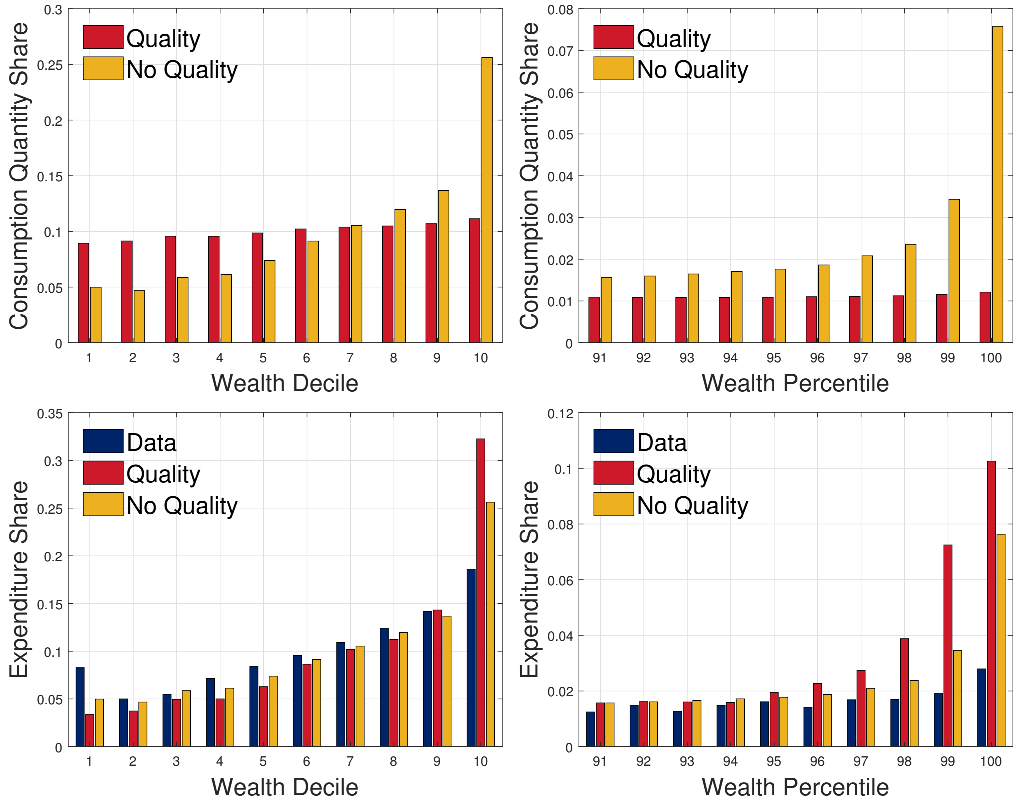

Figure 7: Consumption Expenditure Shares by Wealth – Model vs. Data

Notes: The figure shows consumption shares (top panels) and expenditure shares (bottom panels) by wealth deciles and percentiles in the calibrated models with and without quality. Expenditure data are from the 2019 PSID wave for household heads of 25-65 years old. Wealth is total assets minus total liabilities at the household level. See Online Appendix A for details on variables' definitions and construction.

Second, the two calibrated models have remarkably similar implications for consumption expenditures, although for radically different reasons. In the standard model, expenditures move one-for-one with consumption quantity. In the model with quality, instead, expenditures closely track unit prices, while consumption quantity changes little across households. An immediate implication of these results is that one cannot readily use expenditure data to discriminate between the two models. Yet, one can still use expenditure data to assess the extent to which the expenditure patterns in the model with quality are borne out in the data and whether expenditure shares in the data vary across good categories in a way that can be associated with consumption quality.

To be sure, mapping the consumption good in the model to a specific good category in the data is problematic and, to a large extent, unwarranted. In addition, quality measurement remains an empirical challenge, even if one has access to information on unit prices. To bypass these issues, for the moment, we proceed to examine expenditure patterns for luxury and necessity goods in PSID. The underlying idea is that luxuries are goods for which a broader spectrum of quality levels are available than necessity goods. And to the extent that this is the case, one would expect the distribution of expenditures on luxuries to be highly dispersed and more concentrated at the top of the wealth distribution relative to expenditures on necessities. Reassuringly, we find that this conjecture stands true in the data.

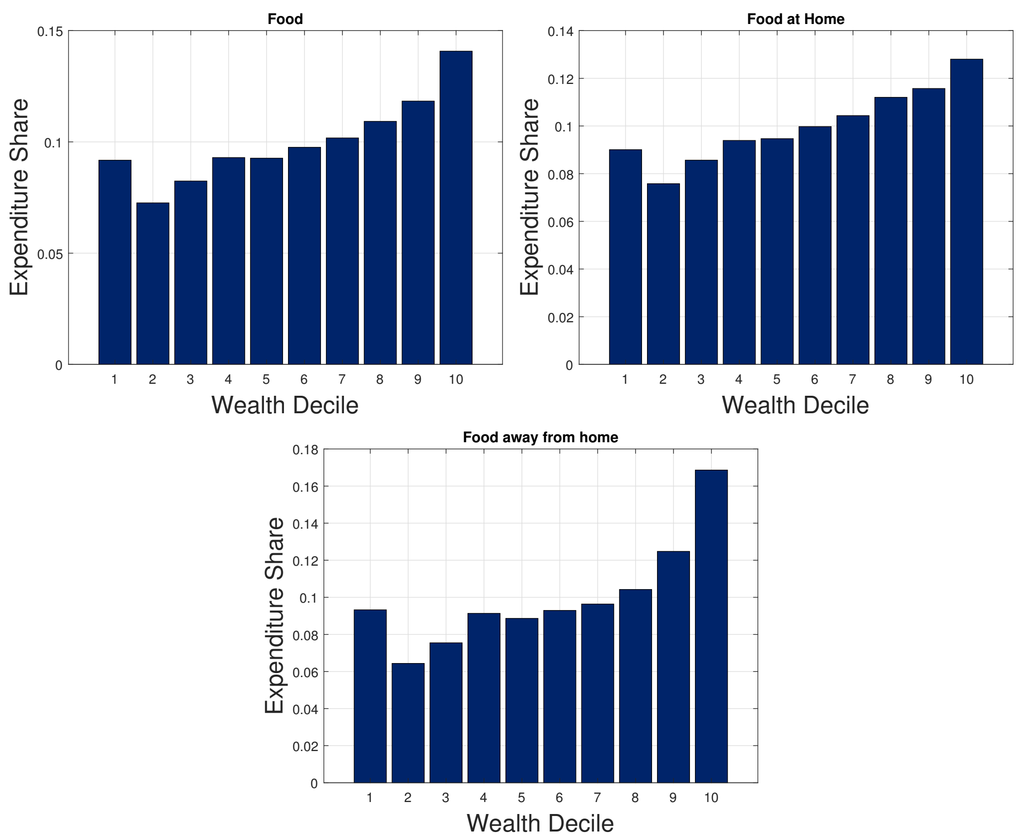

Figure 8: Food Expenditure Shares in the Data

Notes: The figure shows the distribution of total food expenditure (top left panel), expenditure on food at home (top right panel), and expenditure on food away from home (bottom panel), by wealth deciles. Data are from the biannual 2005-2015 PSID waves for households heads of 25-65 years old. See Online Appendix A for details on variables' definitions and construction.

Consumption expenditures in PSID. We consider PSID data on expenditures per person on food at home, food away from home, clothing, entertainment, education, and vacation.[12] The literature typically associates food at home with a "necessity," and the other five categories with "luxuries" (see Aguiar and Bils, 2015; Chang, Hornstein and Karabarbounis, 2020). We find that household-level data from the PSID provides large support to the predictions of the model with the quality choice. In the model, as in the data, consumption expenditures are unevenly distributed across the wealth distribution.

We consider food expenditures first. Food is an appealing goods category for at least two reasons. First, it is an essential good required to sustain life, implying that it must belong to the consumption basket of all households. Second, given its physical nature, it would seem natural to think that differences in the *quantity* consumed by households with different incomes or wealth are plausibly "small," relative to the quality of the food consumed. The advantage of the the food category is that we can view differences in food expenditures across households as mainly coming from differences in the prices paid.

Figure 8 shows the shares of total food expenditure (top left panel), food expenditure at home (top right panel), and food expenditure away from home (bottom panel) by wealth deciles. (To be precise, we calculate shares relative to the total expenditures on food, food at home, and food away from home, respectively.) Expenditure on food away from home – a "luxury" – is highly concentrated. Households in the top decile of the wealth distribution account for roughly 17% of total expenditure on food away from home, whereas households in the bottom decile account for 9%.[13] In contrast, expenditure on food at home – a "necessity" – is more evenly distributed across the wealth distribution. The first wealth decile accounts for slightly more than 8% of total expenditure on food at home, whereas the last wealth decile accounts for 13%. Hence, while there is variation in expenditure shares, it is not nearly as large as that for food away from home.

Figure 9 shows shares of expenditures on clothing, entertainment, education, and vacation by wealth deciles. Similarly to food away from home, expenditure shares rise with wealth. Indeed, the inequality in consumption expenditures for these luxuries is more pronounced than for food away from home. For example, the top decile of the wealth

[12] For broad goods categories, PSID expenditure data are largely consistent with CEX data (see, e.g., Li et al., 2010; Andreski et al., 2014).

[13] In Online Appendix A, Figure A.13 shows data on food expenditure shares on food away from home by earnings decile. The higher the earnings, the higher the share of food expenditures that goes to food away from home. The dispersion is sizable, going from nearly 27% at the bottom decile to 48% at the top decile.

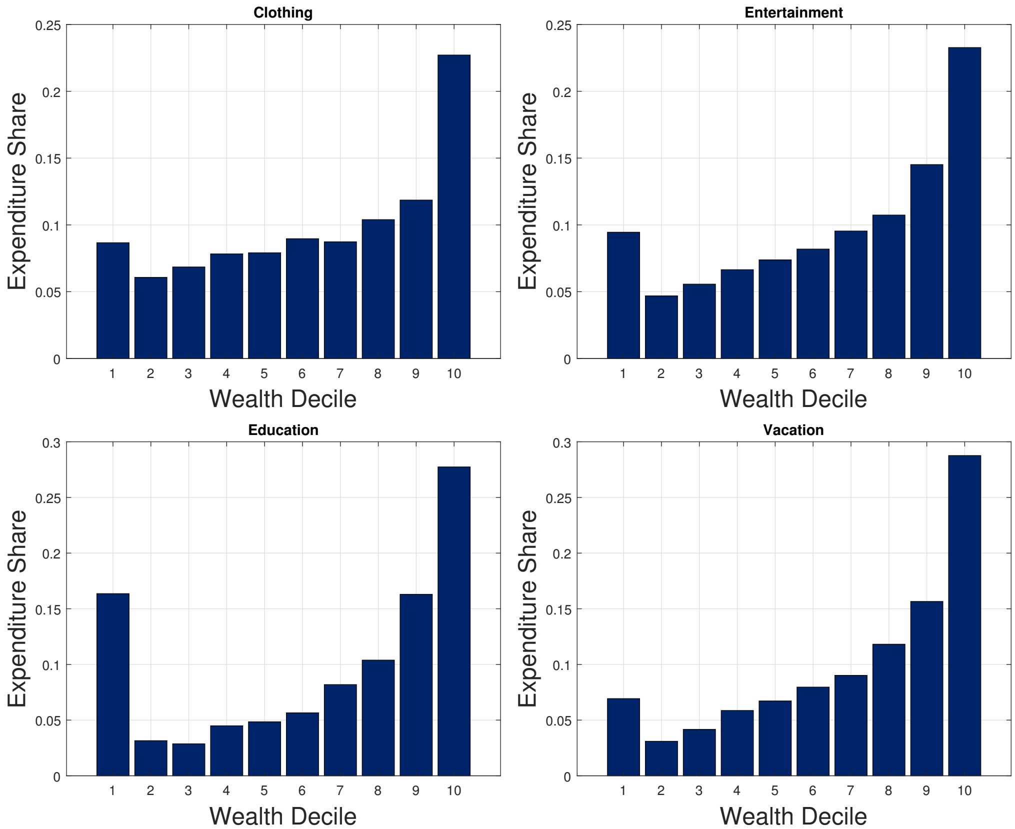

Figure 9: Other Expenditure Shares in the Data

Notes: The figure shows the distribution of expenditure shares on clothing (top left panel), expenditure shares on entertainment (top right panel), expenditure shares on education (bottom left panel), and expenditure shares on vacation (bottom right panel) by wealth deciles. Data are from the biannual 2005-2015 PSID waves for households heads of 25-65 years old. See Online Appendix A for details on variables' definitions and construction.

distribution accounts for nearly 23% of total expenditures on clothing, as opposed to the 17% figure for food away from home. Similar patterns hold for entertainment, education, and vacation.

7. Micro Evidence on Consumer Quality Choices

This section provides evidence on consumption quantity and quality choices from CEX, Nielsen Consumer Panel, and CSFII. We find that unit prices and available measures of quality of food consumption are sensitive to income, whereas consumption quantity is considerably less so if any. These findings provide strong support for the model with the quality choice. A summary is as follows:

1. **Quality Engel curves.** Using data on durables' expenditures and housing from CEX and scanner data from Nielsen, we estimate quality Engel curves relating unit prices to income. In the data, as in the model, average unit prices increase with income. In addition, using CEX data on clothing prices and quantities, we find that the income elasticity of prices is considerably larger than the income elasticity of quantities. We draw similar conclusions from quantity regressions based on Nielsen.

2. **Food consumption.** Using CSFII data, we find that the quantity of consumption measured by total calories is virtually insensitive to household income. In contrast, regressions based on measures of food content, such as vitamins A, C, and E, calcium, cholesterol, and saturated and unsaturated fats, reveal a positive relationship between food quality and income. Higher-income households consume less saturated fats and cholesterol, typically associated with healthier and higher-quality food consumption.

7.1. Quality Engel Curves

In this subsection, we study the relationship between unit prices and income and how it varies across the income distribution using data from CEX and Nielsen. We partition the income variable in quartiles and define a set of dummies, $\mathbb{1}_{ik}$, which equals one if the household i's income lies in quartile k, and zero otherwise. We consider the following regression:

$$\log\left(p_{ijt}\right) = \alpha + \sum_{k=1}^{4} \beta_k \mathbb{1}_{ik} + \gamma X_{it} + \epsilon_{ijt}, \tag{28}$$

where the subscript i identifies the household, j identifies the retail category in Nielsen, e.g., grocery, or the durable good and housing (rent and rent equivalent) in CEX, p_{ijt} is the average unit price for the regressions based on Nielsen data, and expenditures on durables and housing for the regressions based on CEX data, X_{it} includes demographic variables, such as age, education, marital and employment status, household size and composition, race, ethnicity, occupation, state of residence, and time fixed effects, and ϵ_{ijt} is a residual from the regression. The exact definition of the variables changes depending on whether we use CEX or Nielsen data; also, the number of covariates varies as some variables can be found in one dataset but not in the other.[14]

The coefficients of interest are the β_k. We take the first quartile as the reference point so that β_k represent percent differences in average unit prices paid by each income quartile relative to the first income quartile. Loosely speaking, β_k reflects two types of variation in the data. The first type of variation is within the retail category and comes from higher-income households paying, on average, higher unit prices for a given category. The second type of variation is between retail categories and comes from higher-income households with higher expenditure shares in more expensive categories.

Similarly, for the CEX regressions, the relevant cross-sectional variation comes from within and between durables and housing categories. Again, we do not attach any causal interpretation to these estimates. Instead, we use them to gauge the plausibility of the quality Engel curves generated by the model.

Nielsen. The Nielsen Consumer Panel dataset includes data on food and non-food items purchased by a large panel of households across multiple retailers in the United States. The dataset records the actual price paid by the household and the quantity purchased at the Universal Product Code (UPC) level. In addition, it contains detailed demographics about the shopper making the purchases and a categorical variable for total household income. The main advantage of Nielsen over other available datasets, for example, PSID and CEX, is that it allows us to disentangle prices and quantities by product and retail category.[15]

Table 1 reports the estimates based on Nielsen for two specifications of (28) with and without demographic controls. For the specification with demographic controls, we find that households in the top income quartile pay, on average, 24% more per item than households in the lowest income quartile.[16] This figure is broadly consistent with estimates in Jaimovich et al. (2019b) but much smaller than their and our estimates based on durables and housing in the CEX.[17] Table 3 further shows a positive relationship between

[14]See Online Appendix A for details on variables' definitions and construction for the CEX and Nielsen.

[15]Specifically, we calculate the unit price as the total price paid for all units before the discount, minus the total discount due to coupons, and divide this amount by the number of units purchased.

[16]In Online Appendix A, Figure A.14 shows that similar patterns hold within six retail categories, i.e., grocery, discount store, apparel store, convenience store, drug store, and restaurant.

[17]In Online Appendix A, Table A.9 reports estimates from quantity regressions based on Nielsen. Overall, we find that households in the top income quartile consume approximately 12% less than households in the first income quartile. Figure A.15 shows the regression coefficients on the household income dummies, β_k, for six retail categories, i.e., grocery, discount store, apparel store, convenience store, drug store, and restaurant. We find substantial heterogeneity in β_k across categories in terms of both signs and income profiles. For example, for grocery, β_k are positive and decreasing in income. In contrast, for discount

unit prices and employment.

Table 1: Unit Prices and Income – Nielsen

	log (Price)	log (Price)
Relative to income quartile 1		
Income quartile 2	0.113***	0.095***
	(0.003)	(0.004)
Income quartile 3	0.186***	0.162***
	(0.003)	(0.004)
Income quartile 4	0.270***	0.241***
	(0.004)	(0.005)
Demographic controls	✗	✓
Time fixed effects	✓	✓
Observations	2,497,867	1,952,202

Notes: The table reports the log-differences in average unit prices paid by each income quartile relative to the first income quartile for household heads of 25-55 years old. Data are from the Nielsen dataset for 2010-2019 (5% random sample). Regressions include a constant. Standard errors clustered at the household level in parentheses. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$. See Online Appendix A for details on variables' definitions and construction.

CEX. Table 2 reports the estimates based on durables and housing expenditure data from CEX. We note that a positive elasticity of expenditures to income is not evidence that higher-income households buy higher-quality goods. A positive relation between expenditures and income could result from higher-income households buying more of the same goods. While valid in principle, such a concern is much less relevant for durables. Durables are well-known to be infrequently purchased and indivisible, so that one can expect a slight variation in the quantity purchased by households at a quarterly frequency.

stores, β_k are negative and increasing in income: households in the fourth income quartile purchase 20% less than households in the first income quartile. Across all specifications and retail categories, we find that the quantity of consumption is far less sensitive to income than unit prices.

For the specifications with demographic controls, we find that households in the top income quartile pay approximately 60% more than households in the lowest income quartile. This figure is again consistent with available estimates in the literature (see, e.g., Bils and Klenow, 2001; Jaimovich et al., 2019b), and supports the view that quality Engel curves are steeper for durables for which there is arguably a broader array of quality-price bundles from which to choose.[18] Similarly to Nielsen, Table 3 confirms a positive relationship between durables and housing expenditures and employment in CEX.

7.2. Food Consumption

The appealing feature of the CSFII that makes it uniquely suitable for measuring quality choices is that it contains measures of *food intake* at the individual level using detailed food diaries, including the quantity and the quality of food consumption.[19] We focus on prime-age household heads 25-55 years old. The two waves of the CSFII include diaries from 1989-91 and 1994-96, which we pool as a single cross-section and include year dummies in the regressions.

Our estimating regression model is

$$\log(intake_{it}) = \alpha + \beta \log(income_{it}) + \gamma X_{it} + \epsilon_{it}, \tag{29}$$

where the dependent variable $intake_{it}$ is calories, vitamin A, C, E, calcium, cholesterol, saturated and unsaturated fats, and proteins, all measured in grams, for household head i in survey year t, $income_{it}$ is total household income, and X_{it} is a vector of covariates that includes standard demographic characteristics (age, gender, and race), household size, and dummies for survey years, region and metropolitan area of residence, height, and a large number of health variables.

Table 4 reports estimates of the income elasticity of food intake from ordinary least squares (OLS) and instrumental variable (IV) regressions. To deal with the problem of measurement error in income and unmeasured omitted variables, we follow Aguiar and Hurst (2005) and instrument household income with occupation, education, education and occupation interactions, and gender and race interactions. Aside from the log-calories regression, all other regressions include log calories as an additional control.

[18] In Online Appendix A, using CEX data on clothing prices and quantities, we show in Table A.8 that the income elasticity of prices is considerably larger than the income elasticity of quantities.

[19] The CSFII has also been used by early studies of food expenditures over the life cycle (Aguiar and Hurst, 2005, 2013).

Table 2: Unit Prices and Income – CEX

	Durables		Housing	
	log (Price)	log (Price)	log (Price)	log (Price)
Relative to income quartile 1				
Income quartile 2	0.229***	0.225***	0.429***	0.260***
	(0.009)	(0.010)	(0.005)	(0.005)
Income quartile 3	0.380***	0.398***	0.666***	0.414***
	(0.009)	(0.011)	(0.006)	(0.007)
Income quartile 4	0.648***	0.677***	0.961***	0.607***
	(0.009)	(0.013)	(0.010)	(0.010)
Demographic controls	✗	✓	✗	✓
Time fixed effects	✓	✓	✓	✓
Observations	630,164	544,599	88,328	71,658

Notes: The table reports the log-differences in average unit prices paid by each income quartile relative to the first income quartile estimated based on CEX data for 2004-2019. Durables include: new motor vehicles, new vehicle accessories, used vehicles, furniture, glassware, other equipment, hardware/tools, televisions, audio/video equipment, computers and accessories, video games, recording media, sporting equipment, supplies, guns and ammunition, toys games and hobbies, bicycles and accessories, pleasure boats, other recreational vehicles, recreational books, other books, musical instruments, jewelry and watches, telephone and fascimile equipment, medical equipment, property, capital improvement materials, other electronics, luggage, miscellaneous durables. Housing includes rent and rent equivalent. Regressions include a constant. Robust standard errors in parentheses. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$. See Online Appendix A for details on variables' definitions and construction.

Also, for the log fat regressions, the log of total fats is included as an additional control.

Calories vary slightly with income within the cross-section of prime-age individuals. However, other food intake components are strongly correlated with income. Specifically, the income elasticities of vitamins and polyunsaturated fat, a "good fat," are positive and statistically significant at the 1% level in both OLS and IV regressions. In contrast, income elasticities of cholesterol and saturated fat, "a bad fat," are negative but similarly

Table 3: Unit Prices and Employment

	CEX		Nielsen
	Durables	Housing	
Employment	0.117***	0.156***	0.037***
	(0.013)	(0.009)	(0.005)
Demographic controls	✓	✓	✓
Time fixed effects	✓	✓	✓
Observations	443,381	65,799	1,952,202

Notes: The table reports the coefficients on employment dummies (30+ hours of work) from regressions based on the CEX and Nielsen samples as in Tables 1-2. Regressions include a constant. Standard errors in parentheses. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$. See Online Appendix A for details on variables' definitions and construction for the CEX and Nielsen.

statistically significant at the 1% level.

Overall, the results provide strong evidence that the nutritional quality of food consumption deteriorates at the lower end of the household's income distribution. Individuals consume inexpensive calories by switching their food consumption toward saturated fats and cholesterol and away from vitamins, calcium, and unsaturated fats. These results are consistent with the idea that "healthy diets" are expensive and cannot be afforded by poor households.[20]

8. Implications for Taxes and Transfers

This section assesses the extent to which the consumption quality choice changes the standard model's implications for the effects of changes in taxes and transfers. Before proceeding, note that one could conduct the same exercise in a model with a quality

[20] In Online Appendix A, Tables A.6-A.7 report regression results for food expenditures at home (a "necessity") and away from home (a "luxury") based on CSFII data. We find a positive and statistically significant income elasticity of food expenditure away from home and a negative and statistically significant income elasticity of food expenditure at home.

Table 4: Income Elasticity of Food Intake

	OLS	IV
Calories	-0.000	-0.014
	(0.010)	(0.028)
Vitamin A	0.121^{***}	0.479^{***}
	(0.024)	(0.069)
Vitamin C	0.139^{***}	0.399^{***}
	(0.025)	(0.072)
Vitamin E	0.066^{***}	0.182^{***}
	(0.012)	(0.035)
Calcium	0.016	0.071^{**}
	(0.011)	(0.033)
Cholesterol	-0.071^{***}	-0.217^{***}
	(0.017)	(0.046)
Saturated fat	-0.020^{***}	-0.062^{***}
	(0.005)	(0.014)
Polyunsaturated fat	0.036^{***}	0.148^{***}
	(0.009)	(0.026)
Monounsaturated fat	-0.000	-0.019^{**}
	(0.003)	(0.009)
Protein	0.007	-0.003
	(0.008)	(0.022)

Notes: Data is from the 1989-91 and 1994-96 waves of the CSFII for household heads of 25-55 years old. The table reports the coefficients on the log of income estimated from OLS and IV regressions of the food intake variable (in logs) on the log of income and a list of control variables, that includes age, gender, race, the highest grade of formal schooling completed, household size, and dummies for survey years, region and metropolitan area of residence, height, and health-related variables (weight, HEALTH, DOCTOR1, DOCTOR2, DOCTOR3, DOCTOR4, DOCTOR5, DOCTOR6, and DOCTOR7). For IV regressions, we instrument the log of income with occupation, education, education-occupation interactions, and sex-race interactions. Regressions include a constant. Aside from the log calories regression, all other regressions include log calories as an additional control. For the log fat regressions, the log of total fats is included as an additional control. First-stage F statistics: 57.35 (calories), 56.13 (vitamin A), 55.77 (vitamin C), 55.99 (vitamin E), 56 (calcium), 55.75 (cholesterol), 54.67 (fats), and 55.99 (protein). Robust standard errors in parentheses. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$. See Online Appendix A for details on variables' definitions and construction.

choice and non-homothetic preferences in which markets are complete so that households achieve perfect insurance against shocks. For example, a viable approach would be to feed into a complete-markets version of the model, the wealth distribution from the data, and look at the model implications for hours worked before and after a change in taxes or transfers. However, there are at least two caveats to this approach.

First, extensive literature finds self-insurance as critical to understanding the cross-sectional distributions of consumption, savings, and hours (see, e.g., Blundell, Pistaferri and Preston, 2008; Heathcote, Storesletten and Violante, 2008; Kaplan and Violante, 2010; Jappelli and Pistaferri, 2010; Heathcote, Storesletten and Violante, 2014). Hence, an incomplete-markets model with a realistic precautionary saving motive is the natural framework for counterfactual and policy analysis. Second, in the model's equilibrium, the precautionary saving motive interacts with the consumption quality choice. Quantifying such interaction is of interest per se, let alone its policy implications.

8.1. Taxes and Labor Supply Elasticities

As a first step towards understanding how work incentives change in response to changes in labor taxes, it is helpful to look at how the marginal utility of consumption (MUC) varies by wealth. The MUC is a crucial model object that directly affects work incentives. The lower the MUC, the less valuable in terms of utils the additional unit of consumption from working. In this sense, the lower the MUC, the lesser the work incentives.

As shown in Figure 10, there is a stark difference between the profiles of MUCs by wealth in the two models: MUCs rise with wealth in the model with quality, whereas they fall in the model without quality. The declining profile of MUCs is the culprit for the counterfactual profile of employment rates by wealth in the standard model without quality. In both models, the quantity of consumption increases with wealth. As a result, in the standard model, the MUC must fall by wealth; in the model with the quality choice, instead, the fall in MUC due to the increase in the quantity of consumption is more than offset by the rise in the quality of consumption, as wealthier households purchase higher-quality goods. Such a qualitative difference in how work incentives comove with wealth has important implications for labor supply elasticities.

In the standard labor supply model, the substitution effect as measured by the Frisch elasticity is offset by income or wealth effects, such that the Marshallian elasticity of labor supply is smaller than the Frisch elasticity and, in fact, considerably smaller for plausible parameterizations of the model. For example, an increase in after-tax labor income (or its present discounted value) reduces the labor supply, thus counteracting the increase in the labor supply from the substitution effect. This logic also explains why the standard model fails to reproduce the nearly flat wealth-hours profile as in the data. By contrast, in the model with the quality choice, the income or wealth effect is much attenuated, implying less offsetting of the substitution effect and, thereby, a higher Marshallian elasticity for any given level of wealth. Overall, the quality choice with non-homothetic and non-separable preferences raises labor supply elasticities across-the-board and, thereby, the aggregate earnings elasticity.

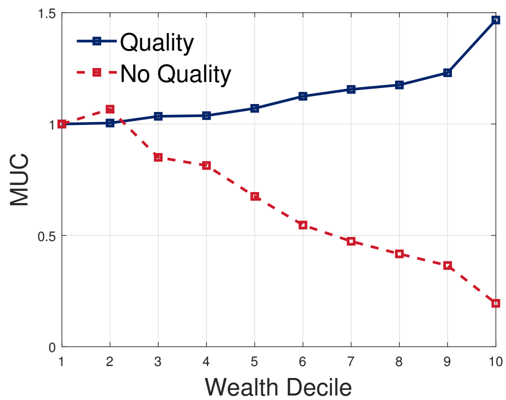

Figure 10: The Marginal Utility of Consumption

Notes: The figure shows the average marginal utility of consumption (MUC) by wealth decile in the calibrated models with and without the quality choice. We divide the MUCs by the MUC in the first wealth decile to facilitate comparison.

To illustrate the differences between the two models' earnings elasticities and as a first attempt to gauge their policy implications, here we calculate optimal tax rates by varying top earnings percentiles for both models with and without the quality choice, using Saez (2001)'s optimal tax formula. Saez derives the optimal tax formula in a static, partial-equilibrium setting in which prices are not allowed to change in response to changes in the tax rate.[21] The optimal top tax rate, $\bar{\tau}$, above a given earnings level $\bar{y}$ depends on four inputs, (1) the average marginal welfare weight, $\bar{g}$, (2) the average uncompensated earnings elasticity, $\bar{\zeta}^u \gtreqless 0$, (3) the average compensated earnings elasticity, $\bar{\zeta}^c = \bar{\zeta}^u - \bar{\eta} \geq 0$, where $\bar{\eta} \leq 0$ subsumes the average income effect, and (4) a tail-statistic $\kappa \equiv y/(y - \bar{y})$, where y in the average of earnings above $\bar{y}$, which measures the thinness of the tail of the earnings distribution above $\bar{y}$,

$$\bar{\tau} = \frac{1 - \bar{g}}{1 - \bar{g} + \bar{\zeta}^u + \bar{\zeta}^c(\kappa - 1)}. \tag{30}$$

[21] Badel and Huggett (2017) show how to extend Saez's formula to a dynamic environment.

In (30), $\bar{g} \equiv \sum_{y^i > \bar{y}} \lambda^i g^i \left(\frac{y^i - \bar{y}}{y - \bar{y}}\right)$ and $g^i \equiv \frac{u_1(c^i, q^i)/p(q^i)}{\sum_{i=1}^{N} \lambda^i [u_1(c^i, q^i)/p(q^i)]}$ is the normalized marginal welfare weight for individual i whose mass in the stationary equilibrium is λ^i, assuming a utilitarian social welfare function.[22]

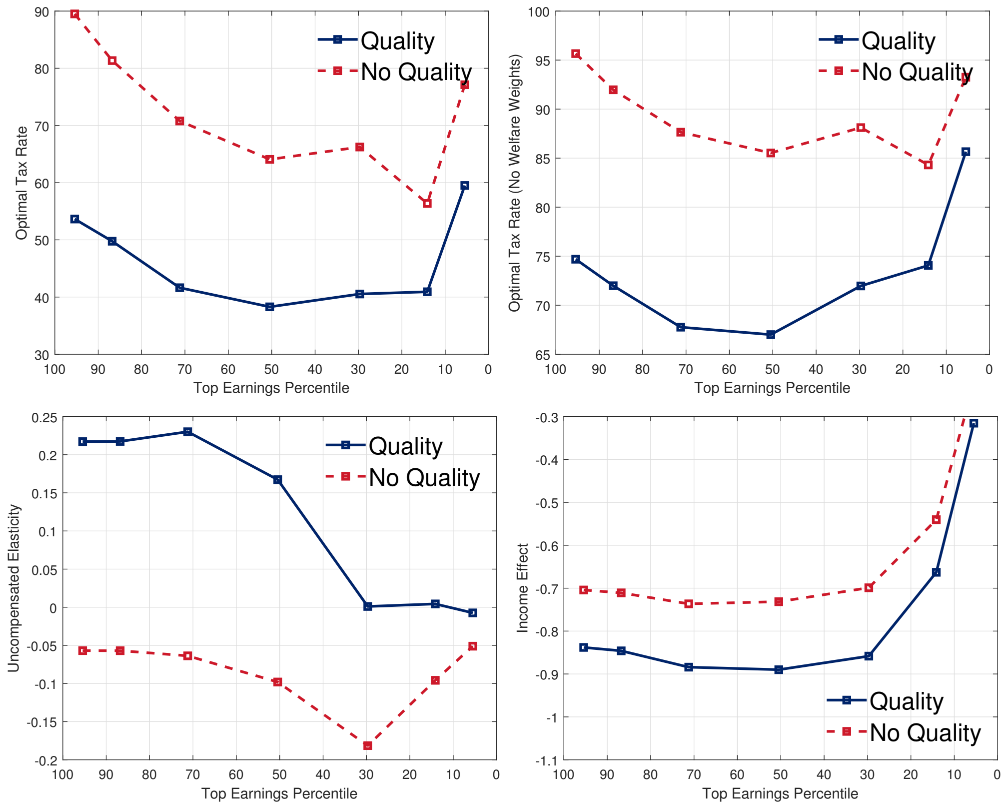

Figure 11: Optimal Tax Rate Simulations

Notes: The figure shows the optimal tax rate, $\bar{\tau}$, by top earnings percentiles calculated based on Saez (2001)'s optimal tax formula reproduced in (30) (top left panel), the optimal tax rates calculated without social welfare weights setting $\bar{g} = 0$ (top right panel), the average (income weighted) uncompensated elasticities $\bar{\zeta}^u$ (bottom left panel), and the average income effects $\bar{\eta}$ (bottom right panel).

Figure 11 shows the optimal tax rates implied by (30) in the calibrated models with and without the quality choice as we increase the value of $\bar{y}$, ranging from the top 95% to the top 5% of the earnings distribution. In both models, optimal tax rates are U-shaped in the threshold earnings level $\bar{y}$. But, importantly, they are considerably higher in the model without the quality choice, a manifestation that the labor supply elasticities in the model with the quality choice are greater than those without the quality choice. As an example, consider the top 5% of the earnings distribution. In that case, the standard model

[22]In the model without the quality choice, consumption quality is fixed at $q = 1$, and the price of the consumption good equals one.

without the quality choice implies an optimal tax rate of 67%, while the model with the quality choice implies a considerably lower one of roughly 60%. The two models also have markedly different implications if one considers the top 95% of the earnings distribution, thus including households typically viewed as poor. In this case, Saez's optimal tax formula prescribes a tax rate as high as 75% in the model without quality compared to the much lower value of almost 55% in the model with quality. In the standard model, low-wealth households, typically low-income, always work and, as such, have very low earnings elasticities. This pattern implies optimal tax rates that are incredibly high at almost expropriation levels at the lower end of the wealth distribution.

Similar patterns hold when we recalculate optimal tax rates without social marginal welfare weights setting $\bar{g} = 0$. As expected, optimal tax rates without welfare weights exceed and, in fact, considerably so the optimal tax rates calculated with welfare weights. The bottom panels of the figure show that the differences in the profile of the optimal tax rates between the two models are due to the differences in the uncompensated labor supply elasticities and income effects as opposed to welfare weights. Such differences in implied earnings elasticities remain large across all the earnings distribution.

8.2. The MPCs Across the Wealth Distribution

In this subsection, we study the implications of quality choice for the marginal propensity to consume (MPC). As customary in the literature, we take a partial-equilibrium approach in which we use the policy functions for the baseline economies with and without the quality choice and compute the responses of consumption expenditures to a one-time change in lump-sum transfers while keeping prices fixed. More specifically, we define the MPC as the fraction of a small, unanticipated transitory transfer (typically around \$500) that a household spends within a given period, say, a quarter.

For a household with assets a and productivity z at the time the transfer of size x is received, the impact MPC is calculated as

$$m(a, z; x) = \frac{p(q(a + x, z))c_q(a + x, z) - p(q(a, z))c_q(a, z)}{x}, \tag{31}$$

where $c_q(a, z)$ and $q(a, z)$ are the policy functions for consumption quantity and quality, respectively, $p(q(a, z))$ is the unit price associated with the quality level $q(a, z)$, and $x =$ \$500, which is approximately the size of common stimulus programs from which MPCs are estimated in the literature (see Jappelli and Pistaferri, 2010, for a survey).

To gauge the differences between the two models, the left panel of Figure 12 shows the ratio of MPCs in the model with quality to the MPCs in the model without quality. We find that the MPCs in the model with quality are generally larger than those in the model without quality: approximately 20% larger across wealth levels, with the exception of the first wealth decile where MPCs are instead similar in magnitude. Importantly, MPCs remain substantially larger in the model with quality even at the top of the wealth distribution, where one would naturally expect the precautionary saving motive to be less important. These results suggest that the quality choice in consumption alters in a significant way the transmission mechanism of policies based on transfers.

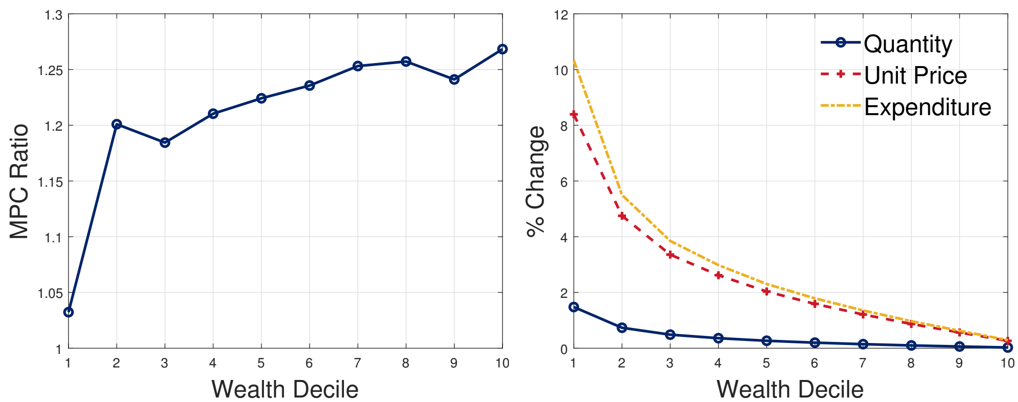

Figure 12: The Marginal Propensity to Consume

Notes: The left panel shows the ratio of the marginal propensities to consume (MPCs) in the model with quality to the MPCs in model without quality by wealth deciles. For a household with assets a and productivity z at the time the transfer of size x is received, the impact MPC is calculated as $m(a,z;x) = \left[p(q(a+x,z))c_q(a+x,z) - p(q(a,z))c_q(a,z)\right]/x$, where $c_q(a,z)$ and $q(a,z)$ are the policy functions for the quantity and the quality of consumption, respectively, $p(q(a,z))$ is the unit price associated with the quality level $q(a,z)$, and $x = \$500$. The right panel shows the percent change in the quantity of consumption, unit prices, and consumption expenditures in the model with the quality choice for an equally-sized transfer of $x = \$500$.

To quantify the mechanisms of the expenditures response to the transfer, the right panel of Figure 12 shows the percent changes in expenditures, consumption quantity, and unit prices by wealth to an equally-sized one-time transfer of $500.[23] In the model with the quality choice, most of the change in consumption expenditures comes from the change in consumption quality. In response to the one-time transfer, households purchase

[23]We calculate the percent change in consumption quantity as $\Delta_c(a,z;x) = \left[c_q(a+x,z)/c_q(a,z) - 1\right] \times 100$, where $x = \$500$, and then average by wealth deciles. Similarly, to calculate the percent changes in expenditures and unit prices, we use the policy function for quality, $q(a,z)$, and the associated unit prices, $p(q(a,z))$.

more expensive, higher-quality versions of the consumption good, leaving consumption quantity, to a large extent, unchanged. While this is true across the board, the effect is more pronounced for households at the higher end of the wealth distribution.

9. Conclusion

We develop a heterogeneous-agent incomplete-markets model with a quality choice in consumption that generates the cross-sectional relation between wealth and employment as in the data. In the United States, employment rates and hours worked are nearly flat across the wealth distribution. Accounting for this fact is a challenge for standard heterogeneous-agent macro models. In these models, wealthier households consume more, enjoy more leisure hours, and work fewer hours. In our model, "quality" is an attribute of the consumption good valued by households, and higher-quality versions of the same good have higher unit prices. With nonhomothetic preferences, consumption quality increases with income and wealth. Furthermore, to the extent that the marginal utility of consumption depends positively on quality, wealthy households may choose to work to afford expensive, high-quality consumption.

To quantify these mechanisms, we calibrate the model and find that it accounts well for the near-zero correlation between wealth and hours worked in U.S. data. Also, in the model and the data, consumption expenditures are uneven across the income and wealth distribution. Such inequality in consumption expenditures does not come from differences in consumption quantity but rather from differences in consumption quality, thus from the higher unit prices paid by wealthier households. The model generates quality Engel curves comparable to those estimated using microdata.

We use the model to evaluate the implications for taxes and transfers. We find that the quality choice with non-homothetic preferences significantly changes the predictions of the incomplete-markets model. For example, employment rates are considerably more sensitive to changes in work incentives relative to the standard incomplete-market model; the income effect on labor supply is generally much attenuated in the model with quality, which implies less offsetting of the substitution effect. Also, most of the consumption expenditure response to changes in taxes or transfers comes from consumption quality, which gives a new mechanism that the literature has so far overlooked.

Altogether, the results in this paper point to the importance of consumption quality in studying salient features of labor allocations at the individual and aggregate levels.

A natural next step is to assess how the cross-sectional distributions of market hours worked vary systematically across countries with different per capita income levels and how those distributions relate to quality choices. While we view these issues as of first-order importance, we leave them for future research.

Data Availability Statement

The data and code underlying this research is available on Zenodo at https://doi.org/10.5281/zenodo.10055981.

References

Aguiar, Mark, and Erik Hurst. 2005. "Consumption versus Expenditure." *Journal of Political Economy*, 113(5): 919–948.

Aguiar, Mark, and Erik Hurst. 2013. "Deconstructing Life Cycle Expenditure." *Journal of Political Economy*, 121(3): 437–492.

Aguiar, Mark, and Mark Bils. 2015. "Has Consumption Inequality Mirrored Income Inequality?" *American Economic Review*, 105(9): 2725–56.

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

Andreski, Patricia, Geng Li, Mehmet Zahid Samancioglu, and Robert Schoeni. 2014. "Estimates of Annual Consumption Expenditures and Its Major Components in the PSID in Comparison to the CE." *American Economic Review*, 104(5): 132–35.

Arnoud, Antoine, Fatih Guvenen, and Tatjana Kleineberg. 2019. "Benchmarking Global Optimizers." *Working Paper*.

Badel, Alejandro, and Mark Huggett. 2017. "The Sufficient Statistic Approach: Predicting the Top of the Laffer Curve." *Journal of Monetary Economics*, 87: 1–12.

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

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

Bick, Alexander, Nicola Fuchs-Schündeln, and David Lagakos. 2018. "How Do Hours Worked Vary with Income? Cross-Country Evidence and Implications." *American Economic Review*, 108(1): 170–99.

Bils, Mark, and Peter J. Klenow. 2001. "Quantifying Quality Growth." *American Economic Review*, 91(4): 1006–1030.

Blundell, Richard, Luigi Pistaferri, and Ian Preston. 2008. "Consumption Inequality and Partial Insurance." *American Economic Review*, 98(5): 1887–1921.

Boppart, Timo, and Per Krusell. 2020. "Labor Supply in the Past, Present, and Future: A Balanced-Growth Perspective." *Journal of Political Economy*, 128(1): 118–157.

Bridgman, Benjamin, Georg Duernecker, and Berthold Herrendorf. 2018. "Structural Transformation, Marketization, and Household Production around the World." *Journal of Development Economics*, 133: 102–126.

Castañeda, Ana, Javier Díaz-Giménez, and José-Víctor Ríos-Rull. 2003. "Accounting for the U.S. Earnings and Wealth Inequality." *Journal of Political Economy*, 111(4): 818–857.

Chang, Yongsung, Andreas Hornstein, and Marios Karabarbounis. 2020. "Labor-Market Wedge under Engel Curve Utility: Cyclical Substitution between Necessities and Luxuries." *Economic Quarterly*, 106(1): 1–18.

Chang, Yongsung, and Sun-Bin Kim. 2006. "From Individual to Aggregate Labor Supply: A Quantitative Analysis Based on a Heterogeneous Agent Macroeconomy." *International Economic Review*, 47(1): 1–27.

Chang, Yongsung, and Sun-Bin Kim. 2007. "Heterogeneity and Aggregation: Implications for Labor-Market Fluctuations." *American Economic Review*, 97(5): 1939–1956.

Chang, Yongsung, Sun-Bin Kim, Kyooho Kwon, and Richard Rogerson. 2019. "Individual and Aggregate Labor Supply in a Heterogeneous Agent Economy with Intensive and Extensive Margins." *International Economic Review*, 60(1): 3–24.

Chatterjee, Satyajit. 1994. "Transitional Dynamics and the Distribution of Wealth in a Neoclassical Growth Model." *Journal of Public Economics*, 54(1): 97–119.

Cociuba, Simona E., Edward C. Prescott, and Alexander Ueberfeldt. 2018. "US Hours at Work." *Economics Letters*, 169: 87–90.

Coibion, Olivier, Yuriy Gorodnichenko, and Dmitri Koustas. 2021. "Consumption Inequality and the Frequency of Purchases." *American Economic Journal: Macroeconomics*, 13(4): 449–82.

Deaton, Angus. 1988. "Quality, Quantity, and Spatial Variation of Price." *American Economic Review*, 78(3): 418–30.

Deaton, Angus, and John Muellbauer. 1980. *Economics and Consumer Behavior.* Cambridge University Press.

Díaz-Giménez, Javier, Andy Glover, and José-Víctor Ríos-Rull. 2011. "Facts on the Distributions of Earnings,Income, and Wealth in the United States: 2007 Update." *Federal Reserve Bank of Minneapolis Quarterly Review*, 34(1): 2–31.

Faber, Benjamin, and Thibault Fally. 2020. "Firm Heterogeneity in Consumption Baskets: Evidence from Home and Store Scanner Data." *Review of Economic Studies*, forthcoming.

Floden, Martin, and Jesper Lindé. 2001. "Idiosyncratic Risk in the United States and Sweden: Is There a Role for Government Insurance?" *Review of Economic Dynamics*, 4(2): 406–437.

Gomme, Paul, B. Ravikumar, and Peter Rupert. 2011. "The Return to Capital and the Business Cycle." *Review of Economic Dynamics*, 14(2): 262–278.

Guner, Nezih, Remzi Kaygusuz, and Gustavo Ventura. 2014. "Income Taxation of U.S. Households: Facts and Parametric Estimates." *Review of Economic Dynamics*, 17(4): 559–581.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante. 2008. "Insurance and Opportunities: A Welfare Analysis of Labor Market Risk." *Journal of Monetary Economics*, 55(3): 501–525.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante. 2014. "Consumption and Labor Supply with Partial Insurance: An Analytical Framework." *American Economic Review*, 104(7): 1–52.

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

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante. 2020. "Presidential Address 2019: How Should Tax Progressivity Respond to Rising Income Inequality?" *Journal of the European Economic Association*, 18(6): 2715–2754.

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

Imrohoroğlu, Ayşe. 1989. "Cost of Business Cycles with Indivisibilities and Liquidity Constraints." *Journal of Political economy*, 97(6): 1364–1383.

Jaimovich, Nir, Sergio Rebelo, and Arlene Wong. 2019. "Trading Down and the Business Cycle." *Journal of Monetary Economics*, 102: 96–121.

Jaimovich, Nir, Sergio Rebelo, Arlene Wong, and Miao Ben Zhang. 2019b. "Trading Up and the Skill Premium." *NBER Macroeconomics Annual*, 34: 1–20.

Jakobsson, Ulf. 1976. "On the Measurement of the Degree of Progression." *Journal of Public Economics*, 5(1-2): 161–168.

Jappelli, Tullio, and Luigi Pistaferri. 2010. "The Consumption Response to Income Changes." *Annu. Rev. Econ.*, 2(1): 479–506.

Kakwani, Nanak C. 1977. "Measurement of Tax Progressivity: An International Comparison." *Economic Journal*, 87(345): 71–80.

Kaplan, Greg, and Giovanni L. Violante. 2010. "How Much Consumption Insurance beyond Self-Insurance?" *American Economic Journal: Macroeconomics*, 2(4): 53–87.

King, Robert G., Charles I. Plosser, and Sergio T. Rebelo. 1988. "Production, Growth and Business Cycles: I. The Basic Neoclassical Model." *Journal of Monetary Economics*, 21(2-3): 195–232.

Koh, Dongya, Raül Santaeulàlia-Llopis, and Yu Zheng. 2020. "Labor Share Decline and Intellectual Property Products Capital." *Econometrica*, 88(6): 2609–2628.

Kuhn, Moritz, and José-Víctor Ríos-Rull. 2016. "2013 Update on the U.S. Earnings, Income, and Wealth Distributional Facts: A View from Macroeconomics." *Federal Reserve Bank of Minneapolis Quarterly Review*, 37(1): 2–73.

Kuhn, Moritz, Moritz Schularick, and Ulrike I Steins. 2020. "Income and Wealth Inequality in America, 1949–2016." *Journal of Political Economy*, 128(9): 3469–3519.

Li, Geng, Robert F. Schoeni, Sheldon Danziger, and Kerwin Kofi Charles. 2010. "New Expenditure Data in the PSID: Comparisons with the CE." *Monthly Labor Review*, 133(2): 29–39.

McGrattan, Ellen R., and Edward C. Prescott. 2003. "Average Debt and Equity Returns: Puzzling?" *American Economic Review*, 93(2): 392–397.

Mendoza, Enrique G., Assaf Razin, and Linda L. Tesar. 1994. "Effective Tax Rates in Macroeconomics: Cross-Country Estimates of Tax Rates on Factor Incomes and Consumption." *Journal of Monetary Economics*, 34(3): 297–323.

Musgrave, Richard A. 1959. *The Theory of Public Finance: A Study in Public Economy.* New York: McGraw-Hill.

Mustre-del-Río, José. 2015. "Wealth and Labor Supply Heterogeneity." *Review of Economic Dynamics*, 18(3): 619–634.

Oberfield, Ezra, and Devesh Raval. 2021. "Micro Data and Macro Technology." *Econometrica*, 89(2): 703–732.

Pijoan-Mas, Josep. 2006. "Precautionary Savings or Working Longer Hours?" *Review of Economic Dynamics*, 9(2): 326–352.

Restuccia, Diego, and Guillaume Vandenbroucke. 2013. "A Century of Human Capital and Hours." *Economic Inquiry*, 51(3): 1849–1866.

Rogerson, Richard. 1988. "Indivisible Labor, Lotteries and Equilibrium." *Journal of Monetary Economics*, 21(1): 3–16.

Saez, Emmanuel. 2001. "Using Elasticities to Derive Optimal Income Tax Rates." *Review of Economic Studies*, 68(1): 205–229.

Smith, Matthew, Danny Yagan, Owen Zidar, and Eric Zwick. 2019. "Capitalists in the Twenty-First Century." *Quarterly Journal of Economics*, 134(4): 1675–1745.

Yum, Minchul. 2018. "On the Distribution of Wealth and Employment." *Review of Economic Dynamics*, 30: 86–105.