

Tax-and-Transfer Progressivity and Business Cycles*

Youngsoo Jang

Takeki Sunakawa

Minchul Yum

July 2023

Abstract

This paper studies how tax-and-transfer progressivity influences aggregate fluctuations when interacting with household heterogeneity. Using a simple static model of the extensive margin labor supply, we analytically characterize how a degree of progressivity influences differential labor supply responses to aggregate conditions across heterogeneous households. We then build a quantitative dynamic general equilibrium model with both idiosyncratic and aggregate productivity shocks and show that it delivers moderately procyclical average labor productivity and a large cyclical volatility of aggregate hours relative to output. Our quantitative exercises suggest that progressivity at the bottom of the income distribution shaped by the phasing out of transfers is key for these findings. Finally, we provide suggestive empirical evidence on the heterogeneity of employment responses across the wage distribution.

Keywords: Progressivity, targeted transfers, extensive margin labor supply, business cycles, redistributive policies

JEL codes: E32, E24, H31, H53, E21

*Jang: University of Queensland, Australia, e-mail: youngsoo.jang@uq.edu.au. Sunakawa: Hitotsubashi University, Japan, e-mail: takeki.sunakawa@gmail.com. Yum: University of Southampton and CEPR, United Kingdom, e-mail: m.yum@soton.ac.uk. We thank the editor, Kjetil Storesletten, and three anonymous referees for their helpful comments and suggestions that have improved the paper substantially. We are also grateful to Fabian Kindermann, Alex Ludwig, Eunseong Ma and Yikai Wang for their constructive discussions at various stages of this project, and to Árpád Ábrahám, Björn Brügemann, Yongsung Chang, Jay Hong, Pantelis Kazaris, Aubhik Khan, Sun-Bin Kim, Tom Krebs, Dirk Krueger, Per Krusell, Claudio Michelacci, Matthias Meier, Benjamin Moll, Fabrizio Perri, Shuhei Takahashi, Michèle Tertilt, Julia Thomas, and seminar participants at the Bank of Canada, the Bank of Japan, EUI, Federal Reserve Board, Hitotsubashi, KAIST, Korea Univ., Ohio State Univ., Queen Mary Univ. of London, Sogang, UCL, Univ. of Edinburgh, Univ. of Kent, Univ. of Konstanz, Univ. of Mannheim, Univ. of Queensland, Yonsei, and participants at many conferences for comments. Financial support from the German Research Foundation (DFG) through CRC TR 224 (Project A3) and Japan Society for the Promotion of Science through Grant-in-Aid for Scientific Research (KAKENHI) for Young Scientists B (Project 15K20918) is gratefully acknowledged. A previous version of this paper was circulated with the title "Is Household Heterogeneity Important for Business Cycles?".

1 Introduction

A large body of literature has investigated the macroeconomic implications of the progressive nature of taxes and transfers.¹ A natural yet relatively unexplored question is how the progressivity of taxes and transfers affects aggregate fluctuations. This paper asks how the degree of tax-and-transfer progressivity alters the way aggregate shocks are transmitted to the macroeconomy with endogenous labor supply at the extensive margin.²

We begin by considering a simple static model of the extensive margin labor supply, with agents who differ in their potential earnings (high or low) and assets according to a distribution featuring a high concentration near zero. We consider higher progressivity a variation in the transfer schedule, such that agents with low potential earnings (the low type) receive more than those with high potential earnings (the high type). We show that higher progressivity induces the labor supply of the low type to respond more strongly to the aggregate shifter. This is because the threshold asset relevant to the employment decision moves closer to zero, around which there is a higher density of (marginal) agents. Consequently, higher progressivity leads to a lower cyclical variation of average labor productivity through changes in the composition of workers (Bils 1985), and potentially to a greater volatility of aggregate hours driven by the low type.

To explore the role of tax-and-transfer progressivity quantitatively, we then consider a standard incomplete-markets framework with heterogeneous households who make consumption-savings and extensive margin labor supply decisions in the presence of both idiosyncratic productivity risk and aggregate risk (Chang and Kim 2006). Since various welfare programs for low-income households are phased out as income rises, overall tax-and-transfer progressivity may vary over the income distribution with it being particularly high at the bottom of the income distribution.³ To better replicate this empirical pattern, we use two parsimonious yet flexible nonlinear functions for taxes and transfers, separately. We calibrate our model economy to match some salient features in the micro-level data, including the degree of progressivity in welfare programs in the Survey of Income and Program Participation (SIPP) data.

We find that our baseline model delivers aggregate labor market dynamics that differ con-

¹The research questions that have been considered include normative ones (such as optimal progressivity) and positive ones (such as the role of progressivity in explaining macroeconomic outcomes). For example, see Conesa et al. (2009), Heathcote et al. (2014), Bick and Fuchs-Schündeln (2018), and Guner et al. (2020) among others.

²The sizes of welfare programs in the United States—e.g., cash transfers, and food, medical, and childcare support to low-income households—have been steadily growing since the 1970s (Ben-Shalom et al. 2011). Nonetheless, some recent papers found that progressivity has had no clear trend during the post war period, despite a few drastic ups and downs (e.g., Ferriere and Navarro (2022) who focused on taxes, and Heathcote et al. (2020) who considered both taxes and transfers among working-age individuals).

³For example, Fleck et al. (2021) and Ferriere et al. (2022) find that a single log-linear progressive taxation function is not sufficient to approximate the observed tax and transfer scheme, especially at the bottom of the income distribution where progressivity is disproportionately larger.

siderably from its nested versions that abstract either from transfers entirely, from differences in transfers across households, or from household heterogeneity. Specifically, it generates considerably lower correlations between average labor productivity and output (0.69) compared to all the nested models (between 0.84 and 0.95, as compared to 0.30 in the data). At the same time, the cyclical volatility of aggregate hours relative to output is 0.73 in the baseline model, which is higher than 0.51 and 0.60 in the nested heterogeneous-agent models and is much closer to 0.80 obtained from its representative-agent counterpart: a version of a Hansen–Rogerson economy that is known to be successful in generating a high volatility of hours.⁴

It is important to note that the model specifications with household heterogeneity do not achieve the quantitative success in the absence of transfers, even though we maintain the degree of tax progressivity. This suggests that transfers play important roles in delivering the above results. First, with the transfer schedule entered separately that enables us to better replicate progressivity at the bottom of the income distribution, our baseline quantitative model can deliver the theoretic mechanism highlighted in the analytical model more effectively. We visualize this by using impulse response functions at the disaggregated level that show that low productivity households are more responsive to changes in aggregate shocks in our baseline model with progressive transfers, when compared to one without them. The second role arises due to risk and market incompleteness. In the absence of any transfers, the labor supply of low productivity households is highly inelastic irrespective of aggregate conditions for precautionary reasons. The presence of transfers mitigates this precautionary motive, thereby raising the responsiveness of their labor supply to aggregate shocks.

Next, we conduct another counterfactual exercise where we raise progressivity by using either the tax function or the transfer function. We find that higher progressivity through a faster phase-out of the transfer system further reduces the correlation between average labor productivity and output and raises the volatility of hours. On the other hand, changes in tax progressivity holding the transfer scheme unchanged has limited effects on aggregate labor market dynamics. These results reinforce the above finding that progressivity shaped by the rate at which welfare transfers are phased out matter for labor market fluctuations with labor supply at the extensive margin.

Finally, we use micro data from the Panel Study of Income Dynamics (PSID) to empirically explore the heterogeneity of employment changes. Specifically, we find that the individual-level probability of adjusting labor at the extensive margin is significantly higher among low-wage workers, and that the full-time employment rate has fallen more in lower wage quintiles during

⁴Hagedorn and Manovskii (2008) find that in models with labor search frictions, unemployment benefits closer to the potential wage make the value of working closer to that of being unemployed. This in turn increases labor market volatilities, and we note similarities with our volatility results in this regard. However, our key model mechanism relies on *heterogeneity* across households in terms of wealth and the size of transfers. This enables us to go beyond labor market volatilities and study issues related to the cyclical volatility of average labor productivity.

the most recent recessions.

Although extensive studies have shown the importance of heterogeneity in accounting for macroeconomic aggregates and equilibrium prices in the absence of aggregate risk (Huggett 1993; Heathcote 2005), the earlier literature that considers aggregate uncertainty often found that incorporating micro-level heterogeneity has only limited impacts on the business cycle fluctuations of macroeconomic aggregates (e.g., Krusell and Smith 1998; Khan and Thomas 2008; Chang and Kim 2014). Our main result—that household heterogeneity at the micro level can be important for understanding the dynamics of macroeconomic variables—is broadly in line with recent papers, such as Krueger et al. (2016), and Ahn et al. (2017).

Weak correlations between average labor productivity and output or hours—often referred to as the Dunlop–Tarshis observation—are known to be difficult to explain using standard real business cycle models. The literature has suggested various mechanisms to dampen strongly positive correlations, with earlier studies relying on the introduction of additional shocks to representative-agent models such as home-production technology shocks (Benhabib et al. 1991), government spending shocks (Christiano and Eichenbaum 1992), and income tax shocks (Braun 1994). Recently, Takahashi (2020) reduces the correlation between average labor productivity and hours by incorporating uncertainty shocks into a standard heterogeneous-agent model (Chang and Kim 2007). Our result is distinct from the existing literature because our mechanism relies on the existence of institutional features leading to heterogeneous responses.

Our quantitative model highlights the effect government transfers have on the precautionary behavior of poor households. Hubbard et al. (1995) show that social insurance discourages precautionary savings among low-income households. Using an incomplete-markets model without aggregate uncertainty, Yum (2018) finds that transfers are important in bringing the employment rate of wealth-poor households closer to the data, which has important implications for the long-run employment effects of labor taxes. Our results suggest that they have important implications for the dynamics of macroeconomic aggregates over the business cycle as well.

The rest of this paper is organized as follows. Section 2 presents the simple static model and present the analytic results on its key mechanism. Section 3 introduces the quantitative dynamic models. Section 4 explains calibration and shows the properties of the quantitative models in stationary equilibrium. Section 5 presents the main quantitative results. Section 6 presents empirical supporting evidence. Section 7 then presents our conclusions.

2 A static model of the extensive margin labor supply

In this section, we present a simple static model of the extensive margin labor supply.⁵ The goal of this section is to illustrate the *direct* effects of fiscal policy on aggregate labor market fluctuations in a tractable way. This tractability is achieved by the simplifying assumption that the distribution of wealth is fixed with respect to changes in fiscal policy, and that it is independent of potential earnings. As fiscal policy may change the distribution of wealth, and because this *indirect* effect could potentially affect the theoretical predictions of this section, we will explore this mechanism using a more realistic dynamic model in the subsequent sections.

Our model here considers a continuum of agents in the unit interval. We assume that there are two types of agents with different potential earnings (or wage offers). That is, the individual component of the potential wage can be either low or high: $x_i \in \{x_l, x_h\}$. The mass of each type is denoted by π_l and π_h satisfying the condition $\pi_l + \pi_h = 1$. Agents also differ in their level of asset holdings a_i , and they can choose to either work full-time or not at all: $n_i \in \{0, 1\}$.

The decision problem of each type i is given by:

$$\max_{c_i \geq 0, n_i \in \{0, 1\}} \{\log c_i - bn_i\}$$

subject to

$$c_i \leq zx_i n_i + a_i + T_i,$$

where c denotes consumption and $b > 0$ captures the disutility of work. We set $b = \log(2) > 0$ without loss of generality. We use z to denote an aggregate shifter of potential earnings and consider its small perturbations to be the source of aggregate fluctuations. To study the role of progressivity, we allow T_i to depend on the type i (or potential wages).⁶

By comparing the utility conditional on working to that on not working, the agent chooses to work if:

$$\log(zx_i + T_i + a_i) - b \geq \log(T_i + a_i),$$

⁵Our analytical framework in this section builds on the theoretical framework of Doepke and Tertilt (2016), although the focus of our analysis is different. Whereas their model is based on two gender types and continuous preference heterogeneity, our model is instead based on two wage-offer types and continuous asset heterogeneity. Moreover, our results cover not only labor supply elasticity but also average labor productivity.

⁶We can interpret T_i as transfers *net of taxes*, or equivalently, negative taxes net of transfers. We opt for the former for expositional convenience. In this section, we require T_i to be non-negative. If T_i is negative, this complicates the analytical derivations since consumption could become negative especially for those with low a and low x .

or if

$$a_i \leq zx_i - T_i.$$

This decision rule shows that the agent is more likely to choose to work if the aggregate shifter z or the individual earnings potential x is higher. Also note that the agent is less likely to choose to work if the size of the transfers is higher.

In this model, aggregate employment is determined by both the decision rule and by the distribution. Let $F(a)$ be the conditional (differentiable) distribution function of assets with its probability density being $f(a) = F'(a)$. Specifically, we use the exponential function in our following results. For $a \geq 0$;

$$\begin{aligned} F(a) &= 1 - \exp(-a), \\ f(a) &= F'(a) = \exp(-a). \end{aligned}$$

This density function has a long right tail in its asset distribution, with a large fraction holding low wealth, being in line with the data.

Given the density function and the decision rule, the fraction of agents working (i.e., the employment rate) for each type is given by:

$$N = F(\underline{a}_i) = 1 - \exp(-\underline{a}_i),$$

where

$$\underline{a}_i = zx_i - T_i. \tag{1}$$

In other words, the employment rate N_i of the type i is the integral of all those type i agents whose asset level is lower than the threshold level \underline{a}_i . We now present some theoretical results based on this model, with all proofs provided in Appendix A.

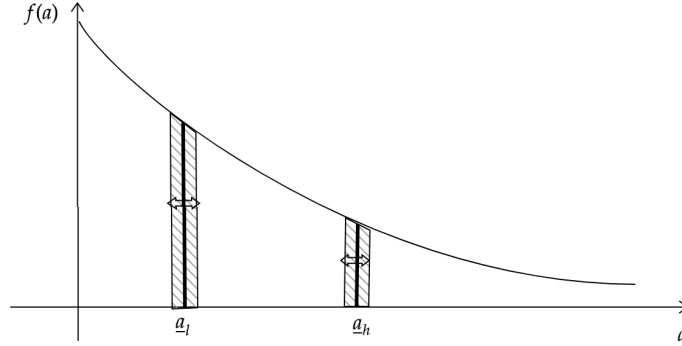
Proposition 1 *Let ε_i be the labor supply elasticity of the type i agents:*

$$\varepsilon_i \equiv \frac{\partial N_i}{\partial z} \frac{z}{N_i}.$$

Assume $T_i = 0$. The labor supply elasticity of agents with low potential earnings is greater than that of agents with high potential earnings. That is, $\varepsilon_l > \varepsilon_h$.

This shows that our model naturally delivers the heterogeneity of labor supply elasticity. The shape of the wealth distribution and the relative location of threshold assets are important for this result. To see this, note that the threshold asset level for the low-potential-wage agents is

Figure 1: Visual illustration of Proposition 1



Note: The shaded areas denote the relevant density of agents who are affected by perturbations in the aggregate shifter z .

lower than that for the high-potential-wage agents: $\underline{a}_l < \underline{a}_h$. As shown in Figure 1, the density of the distribution around \underline{a}_l is greater. Since there are more marginal agents around \underline{a}_l , the same change in the aggregate shifter z —which perturbs both \underline{a}_l and \underline{a}_h —will more strongly affect the employment rate of the low-potential-wage agents.

We now consider the role of government transfers and how they interact with heterogeneity. To simplify the algebra, we impose symmetry. Specifically, we assume that $\pi_l = \pi_h = 0.5$. In addition, $x_h = 1 + \lambda$ and $x_l = 1 - \lambda$, where $\lambda \in (0, 1)$ measures the cross-sectional dispersion.

To study the effects of progressivity, T_i is assumed to be:

$$\begin{aligned} T_l &= T(1 + \omega\lambda), \\ T_h &= T(1 - \omega\lambda), \end{aligned}$$

where $T \in [0, z(1 - \lambda)/2]$ captures the scale of transfers and $\omega \in [0, 1/\lambda]$ shapes the degree of progressivity.⁷ Note that a change in ω does not affect the aggregate size of the transfers.⁸ Given the above assumptions, the employment rates for each type are given by $N_i = 1 - \exp(-\underline{a}_i)$, where $\underline{a}_l = z(1 - \lambda) - T - T\omega\lambda$ and $\underline{a}_h = z(1 + \lambda) - T + T\omega\lambda$.

Proposition 2 *Greater progressivity increases the labor supply elasticity of the low-potential-wage agents, while it decreases the labor supply elasticity of the high-potential-wage agents. That is, $\frac{\partial \varepsilon_l}{\partial \omega} > 0$ and $\frac{\partial \varepsilon_h}{\partial \omega} < 0$.*

⁷The maximum values of T and ω ensure that the threshold assets stay non-negative.

⁸Progressivity can increase through two channels: (i) the scale of the transfers, and (ii) the relative size of the transfers received by low-income households. A change in ω is meant to capture the second channel (i.e., controlling the speed of phasing-out of welfare transfers).

Intuitively, greater progressivity (or a higher ω) shifts \underline{a}_l to the left where the distribution is denser. There, the same change in the aggregate shifter z would induce more agents to change their employment decision, thereby leading to an even larger elasticity for the low-potential-wage agents. By contrast, greater progressivity shifts \underline{a}_h to the right, around which the distribution of assets is thinner. This implies that the elasticity of the high-potential-wage agents should become smaller.

Proposition 3 *Let N denote the aggregate employment rate: $N = \pi_l N_l + \pi_h N_h$. Let ε be the aggregate labor supply elasticity:*

$$\varepsilon \equiv \frac{\partial N}{\partial z} \frac{z}{N}.$$

The aggregate labor supply elasticity is higher with greater progressivity. That is, $\frac{\partial \varepsilon}{\partial \omega} > 0$.

The key to this result is that an increase in the elasticity of low-potential-wage agents should be large enough to outweigh the opposing effects from a decrease in the elasticity of high-potential-wage types. Given that the density function declines at an increasing rate, this condition is satisfied. This result suggests that progressivity potentially has a role in generating large volatility in aggregate hours, as observed in the data.

Finally, we consider the implications for the cyclicity of average labor productivity. We define average labor productivity as output divided by aggregate hours:

$$\chi \equiv \frac{\sum_{j \in \{l, h\}} \pi_j (z x_j N_j)}{\sum_{j \in \{l, h\}} \pi_j N_j} = z \frac{\sum_{j \in \{l, h\}} \pi_j (x_j N_j)}{\sum_{j \in \{l, h\}} \pi_j N_j} \equiv z \chi_0,$$

where we separately define the second term as χ_0 . Here, we can clearly see that a change in the aggregate shifter z would directly cause average labor productivity to become procyclical through the first term z , as is the case in real business cycle models. The second term χ_0 captures the effects through worker composition, which indirectly depends on z through heterogeneous employment responses. The following two propositions focus on this second term.

Proposition 4 *A change in the aggregate shifter z has a direct and an indirect effect on average labor productivity $z \chi_0(z)$. The indirect effect is negative: $\frac{\partial \chi_0(z)}{\partial z} < 0$.*

Proposition 5 *Average labor productivity becomes less positively (or more negatively) correlated with z as progressivity increases: $\frac{\partial}{\partial \omega} \left(\frac{\partial \chi_0}{\partial z} \right) < 0$.*

Proposition 5 shows that progressivity can shape the cyclicity of average labor productivity through worker composition effects. To see what this means, let us suppose that the aggregate

shifter z increases (i.e., in a boom). While both types of agents are more likely to work, relatively more low-type workers would do so when progressivity is greater. This follows from the disproportionate rise in low-type labor supply elasticity shown in Proposition 3. This force would cause lower increases in average labor productivity during booms, thereby dampening the tight positive link between z and average labor productivity.

3 Quantitative business cycle models

As noted earlier, the key results in Section 2 capture the direct effects of fiscal policy changes since they are derived in a static environment. Therefore, it is a quantitative question whether this mechanism would be relevant in a more realistic and dynamic model environment. In the remaining sections, we explore the key mechanisms in models that allow endogenous wealth distributions that can differ by productivity types.

3.1 Baseline model

The baseline quantitative model we use here builds on a standard incomplete-markets framework with both idiosyncratic productivity risk (Huggett 1993; Aiyagari 1994) and aggregate risk, as pioneered by Krusell and Smith (1998). In this model, heterogeneous households make a consumption-savings choice—which endogenizes the distribution of wealth—and a labor supply decision at the extensive margin. There are also differences in transfers across households.

Households The model economy is populated by a continuum of infinitely lived households. It is convenient to describe the decision problem faced by such households in a recursive manner. At the beginning of each period, households are distinguished by their asset holdings a and productivity x_i . We assume that x_i takes a finite number of values N_x and follows a Markov chain with transition probabilities π_{ij}^x from state i to state j . In addition to the individual state variables, a and x_i , there are also aggregate state variables, including the distribution of households $\mu(a, x_i)$ over a and x_i , and aggregate total factor productivity shocks z_k . We also assume that z_k takes a finite number of values N_z following a Markov chain with transition probabilities π_{kl}^z from state k to state l . We assume that these Markov processes of individual productivity x and aggregate total factor productivity (TFP) shock z capture the following continuous AR(1) processes in logarithms:

$$\log x' = \rho_x \log x + \varepsilon'_x, \tag{2}$$

$$\log z' = \rho_z \log z + \varepsilon'_z, \tag{3}$$

where $\varepsilon_x \sim N(0, \sigma_x^2)$ and $\varepsilon_z \sim N(0, \sigma_z^2)$. We denote a variable with a prime symbol its value in the next period. Finally, we assume competitive markets: households take as given the wage rate per efficiency unit of labor $w(\mu, z_k)$ and the real interest rate $r(\mu, z_k)$, both of which depend on the aggregate state variables. Households also take government policies as given.

The dynamic decision problem facing households can then be written as the following functional equation:

$$V(a, x_i, \mu, z_k) = \max \{V^E(a, x_i, \mu, z_k), V^N(a, x_i, \mu, z_k)\},$$

where

$$V^E(a, x_i, \mu, z_k) = \max_{\substack{a' \geq \underline{a}, \\ c \geq 0}} \left\{ \log c - B\bar{n} + \beta \sum_{j=1}^{N_x} \pi_{ij}^x \sum_{l=1}^{N_z} \pi_{kl}^z V(a', x'_j, \mu', z'_l) \right\} \quad (4)$$

subject to

$$\begin{aligned} c + a' &\leq \tau(e, \bar{e})e + (1 + r(\mu, z_k))a + T \\ e &= w(\mu, z_k)x_i\bar{n} \\ T &= T_1 + T_2(m) \\ m &= e + r(\mu, z_k) \max\{a, 0\} \\ \mu' &= \Gamma(\mu, z_k). \end{aligned} \quad (5)$$

and

$$V^N(a, x_i, \mu, z_k) = \max_{\substack{a' \geq \underline{a}, \\ c \geq 0}} \left\{ \log c + \beta \sum_{j=1}^{N_x} \pi_{ij}^x \sum_{l=1}^{N_z} \pi_{kl}^z V(a', x'_j, \mu', z'_l) \right\} \quad (6)$$

subject to

$$\begin{aligned} c + a' &\leq (1 + r(\mu, z_k))a + T \\ T &= T_1 + T_2(m) \\ m &= r(\mu, z_k) \max\{a, 0\} \\ \mu' &= \Gamma(\mu, z_k). \end{aligned} \quad (7)$$

Households maximize utility by choosing their optimal consumption c , their asset holdings in the next period a' , and their labor supply n . Households also face a borrowing limit $\underline{a} \leq 0$. Their labor supply decision is discrete (i.e., $n \in \{0, \bar{n}\}$), and the disutility of work is captured by $B > 0$. Households understand that the expected future value (discounted by a discount factor β) is affected by stochastic processes for individual productivity x' and aggregate TFP

productivity z' , as well as the whole distribution μ' . The budget constraints state that the sum of spending should be less than or equal to the sum of income. The evolution of μ is governed by the law of motion, as denoted by $\mu' = \Gamma(\mu, z_k)$.

As shown in the budget constraints, our model incorporates a progressive tax and transfer system, and these two components are captured separately by two nonlinear functions. First, earnings e are subject to progressive taxation—as is standard in the recent quantitative macroeconomics literature. Specifically, for those who have earnings e , progressive taxation leads to a tax rate of:

$$\tau(e, \bar{e}) = \max \left\{ 1 - \left(\lambda_s (e/\bar{e})^{-\lambda_p} \right), 0 \right\}. \quad (8)$$

Note that, although this function follows the parametric form of Benabou (2002) and Heathcote et al. (2014), we restrict $\tau(e)$ to being non-negative. As is well known, $\lambda_p \geq 0$ captures the degree of progressivity and $\lambda_s \geq 0$ inversely controls the scale of taxation. As the input into the progressive tax schedule is earnings normalized by its average \bar{e} (Guner et al. 2014), a change in λ_p tilts this schedule around average earnings. This strongly affects tax progressivity yet has little effect on the size of taxation.

On top of this typical progressive tax schedule, we also separately introduce progressive transfers. This helps us to better capture progressivity at the bottom of the income distribution, which is difficult to replicate with a single nonlinear tax function (Fleck et al. 2021; Ferriere et al. 2022). Following Krusell and Rios-Rull (1999), we make the specific assumption that transfers T consist of two components. The first component T_1 is given to all households equally, whereas the second component T_2 captures the income security aspect of transfers. In the U.S., there are various means-tested programs, such as the Supplemental Nutrition Assistance Program (SNAP) (formerly known as food stamps), and the Temporary Assistance for Needy Families (formerly the Aid to Families with Dependent Children). As shown in Section 4, the existence of these programs leads us to the observation that the amount of transfers is negatively associated with income. We assume that T_2 depends on total household income m , and use this to replicate the measured transfer progressivity observed in the U.S. data using the following functional form (Yum 2018):

$$T_2(m) = \omega_s (1 + m)^{-\omega_p}. \quad (9)$$

This parametric assumption adds two parameters. First, $\omega_s \geq 0$ is a scale parameter that determines the overall size of the non-flat part of government transfers (i.e., T_2). The next parameter, $\omega_p \geq 0$, governs the degree of progressivity: a higher ω_p makes T_2 decrease faster with income.

Representative firm and government Aggregate output Y is produced by a representative profit-maximizing firm that solves

$$\max_{K,L} \{z_k F(K, L) - (r(\mu, z_k) + \delta)K - w(\mu, z_k)L\} \quad (10)$$

where $F(K, L)$ captures a standard neoclassical production technology in which K denotes aggregate capital, L denotes aggregate efficiency units of labor inputs, and δ is the capital depreciation rate. As is standard in the literature, we assume that the aggregate production function follows a Cobb-Douglas function with constant returns to scale:

$$F(K, L) = K^\alpha L^{1-\alpha}. \quad (11)$$

The first-order conditions for K and L give

$$r(\mu, z_k) = z_k F_1(K, L) - \delta, \quad (12)$$

$$w(\mu, z_k) = z_k F_2(K, L). \quad (13)$$

The government in this economy collects labor taxes from households and uses the tax revenue to finance total transfers to households. The remaining tax revenue is spent as government spending G , which is not valued by households. Note that government spending plays no important role in the exercises of this paper.

Equilibrium A recursive competitive equilibrium is a collection of factor prices $r(\mu, z_k)$ and $w(\mu, z_k)$; household decision rules $g_a(a, x_i, \mu, z_k)$ and $g_n(a, x_i, \mu, z_k)$; government spending G ; a value function $V(a, x_i, \mu, z_k)$; a distribution of households $\mu(a, x_i)$ over the state space; the aggregate capital and labor $K(\mu, z_k)$ and $L(\mu, z_k)$; and the aggregate law of motion $\Gamma(\mu, z_k)$; such that

1. Given factor prices $r(\mu, z_k)$ and $w(\mu, z_k)$, the value function $V(a, x_i, \mu, z_k)$ solves the household decision problems defined above, with the associated household decision rules being:

$$a^* = g_a(a, x_i, \mu, z_k), \quad (14)$$

$$n^* = g_n(a, x_i, \mu, z_k). \quad (15)$$

2. Given factor prices $r(\mu, z_k)$ and $w(\mu, z_k)$, the firm optimally chooses $K(\mu, z_k)$ and $L(\mu, z_k)$ following (12) and (13).

3. Markets clear:

$$K(\mu, z_k) = \sum_{i=1}^{N_x} \int_a ad\mu \quad (16)$$

$$L(\mu, z_k) = \sum_{i=1}^{N_x} \int_a x_i g_n(a, x_i, \mu, z_k) d\mu. \quad (17)$$

4. Government balances its budget. That is, the sum of government spending G and total transfers to households is equal to the total tax revenue.
5. The law of motion for the distribution of households over the state space $\mu' = \Gamma(\mu, z_k)$ is consistent with individual decision rules and the stochastic processes governing x_i and z_k .

3.2 Alternative model specifications

In addition to the baseline model just introduced, we also consider alternative specifications to illustrate the importance of the interplay between household heterogeneity and government transfers.⁹ For convenience, the baseline model featuring "Heterogeneous Agents" and "Targeted" transfers is called Model (HA-T).

The first alternative model specification, denoted as Model (HA-N), is simply a nested specification of the baseline "Heterogeneous-Agent" model with "No" government transfers (i.e., $T_1 = \omega_s = 0$). This model roughly corresponds to the standard incomplete-markets real business cycle model of Chang and Kim (2007), with household heterogeneity and endogenous labor supply at the extensive margin.¹⁰

The second alternative model specification also keeps household heterogeneity but removes differences in transfers across households. We call this model specification Model (HA-F), which is obtained as a nested "Heterogeneous-Agent" model by making transfers "Flat"—that is, independent of income ($\omega_p = 0$). Chang et al. (2013) also consider a business cycle model that is close to this model specification. Note that this form of transfers (flat lump-sum) is very broadly used in the quantitative macroeconomics literature.

Our final alternative specification shuts down household heterogeneity. This "Representative-Agent" version of the model is called Model (RA). Given the indivisible labor supply assumption,

⁹We have also considered a specification which shuts down tax progressivity only. Because its quantitative role is minimal, we have placed those results in Appendix H as a sensitivity check. In Section 5.3, we also consider a counterfactual exercise where we alter tax progressivity using the baseline model specification.

¹⁰A noticeable difference between Model (HA-N) in our paper and the model in Chang and Kim (2007) is that ours includes progressive taxation whereas theirs does not. However, as shown in Section 5 and Appendix H, the business cycle properties of the model are barely affected by the existence of progressive taxation—except for output volatility.

Model (RA) is essentially the business cycle model studied in Hansen (1985) augmented with taxes and transfers. The key feature of this model specification is that the aggregation of Rogerson (1988) under certain assumptions (such as employment lotteries and consumption insurance) leads to the introduction of the stand-in household whose disutility from work is linear—a powerful mechanism used to generate the large volatility of aggregate hours, observed in U.S. data (Hansen 1988). Appendix G includes the detailed model environment and its equilibrium definition.

3.3 Solution method

We solve each of the models numerically. Several key features make the numerical solution method nontrivial for the heterogeneous-agent models. First, the key decision variables in our model are a discrete employment choice and a consumption-savings choice in the presence of a borrowing constraint. Therefore, our solution method is based on a nonlinear method (i.e., the value function iteration) applied to the recursive representation presented above. Second, the aggregate law of motion and the state variables involve an infinite-dimensional object: the distribution μ . This requires us to solve the model by approximating the distribution of wealth as its mean (Krusell and Smith 1998). Since market-clearing is nontrivial in our model with endogenous labor, our solution method also incorporates an additional step to when simulating the model to find market-clearing prices in each period.

We now describe the solution method briefly, with more details found in Appendix G. Following Krusell and Smith (1998), we assume that households use a smaller object that approximates the infinite-dimensional distribution when they forecast the future state variables in order to make current decisions. More precisely, we approximate $\mu(a, x_i)$ by its mean with respect to the asset distribution K . Furthermore, when determining the aggregate capital in the next period K' , real wage per efficiency units w and real interest rate r are assumed to be functions of (K, z_k) instead of (μ, z_k) . We impose parametric assumptions to approximate the aggregate law of motion $K' = \Gamma(K, z_k)$ and $w = w(K, z_k)$ following

$$\hat{K}' = \hat{\Gamma}(K, z_k) = \exp(a_0 + a_1 \log K + a_2 \log z_k) \quad (18)$$

$$\hat{w} = \hat{w}(K, z_k) = \exp(b_0 + b_1 \log K + b_2 \log z_k), \quad (19)$$

as in Chang and Kim (2006, 2007). Households obtain a forecasted \hat{r} based on these forecasting rules, as implied by the first-order conditions of the profit maximization problem facing the representative firm.

The model is solved in two steps. First, given the forecasting rules, we solve for the indi-

vidual policy functions using the value function iterations (*the inner loop*). Then, we update the forecasting rules by simulating the economy using the individual policy functions (*the outer loop*). It is important to once again note that, since our model environment with endogenous labor supply involves non-trivial factor market clearing, we have to incorporate a step to find the market-clearing factor prices in the outer loop (Chang and Kim 2014; Takahashi 2014). We repeat this procedure until the coefficients in the forecasting rules converge.

It is straightforward to solve the representative-agent version of the model. For the purposes of comparison, we keep the same assumptions on the discretization of the TFP shock process as used in the heterogeneous-agent model. The steady-state equilibrium can then be obtained analytically. For solutions with aggregate uncertainty, we use the policy function iteration method.

4 Calibration and model properties in steady state

All model specifications are calibrated to U.S. data. A period in the model is a quarter, as is standard in the business cycle literature. We consider all four of our specifications: Model (HA-T), Model (HA-N), Model (HA-F), and Model (RA).

Calibrating the baseline model We first describe how we calibrate the baseline specification, which involve two sets of parameters. The first set is calibrated externally, in line with the business cycle literature. These parameter values are set in common across all four of our model specifications. The second set of parameters is calibrated to match the same number of relevant target statistics.

We begin by describing the first set of externally calibrated parameters. Most of these are commonly used parameters in the literature. The capital share α is chosen to be consistent with the empirical capital share value of 0.36, and the quarterly depreciate rate δ is set to 2.5%. In our model specifications with a binary labor supply choice, the number of hours worked \bar{n} can be arbitrarily set since it simply determines the scale of the calibrated disutility parameter B . By setting \bar{n} to 1/3—implying that working individuals spend a third of their time endowment on working—we can calibrate $\tilde{B} \equiv B\bar{n}$ directly. Furthermore, the borrowing limit \underline{a} is set to $-T_1/(1+r)$, where r is the equilibrium interest rate in steady state.¹¹

In the literature, tax progressivity λ_p has been estimated using the same functional form we use. As noted by Holter et al. (2019), the estimate of λ_p varies quite a lot (from 0.05 to 0.18), depending on the degree of completeness of the data on government transfers used by researchers.

¹¹This is a form of the natural borrowing limit that ensures that non-working agents are able to pay back their debts in the next period. In Appendix H, we report a version of the model with a zero borrowing limit, as is standard in the literature. The main results found in this paper are robust to this variation.

Because we model progressive transfers separately in addition to progressive taxes, our taxation parameters in (8) should ideally only capture tax progressivity. As the Internal Revenue Service (IRS) income tax data used by Guner et al. (2014) do not include welfare transfers, we use their estimate for $\lambda_p = 0.053$ and $\lambda_s = 0.911$.¹² As discussed below, we then use micro data on the distribution of welfare transfers across households to calibrate the parameters of the transfer function in (9).

The broad goal of this paper is to study how progressive transfers alter the transmission of aggregate shocks in the macroeconomy. As a first step, we consider the most standard one—total factor productivity shocks (Kydland and Prescott 1982)—as an aggregate risk, and employ the standard values of $\rho_z = 0.95$ and $\sigma_z = 0.007$ (Cooley and Prescott 1995). These values are useful as we can easily compare our results to those from recent related papers, such as Chang and Kim (2007) and Takahashi (2020), who also use the same TFP shock estimates.¹³

Next, ρ_x captures the persistence of idiosyncratic risk in the productivity of households. We estimate the persistence of idiosyncratic risk using the PSID following a standard method from the literature (Heathcote et al. 2010), as discussed in Appendix F. The quarterly value based on this estimate is $\rho_x = 0.9847$. The variability of the idiosyncratic risk is calibrated internally and is explained below. We keep the same values for these two parameters, ρ_x and σ_x , for all of the nested model specifications using heterogeneous agents to control for the underlying idiosyncratic risk.

The second set of parameters is jointly calibrated. As shown in Table 1, six parameters are calibrated by matching the same number of target statistics. We now explain how each parameter is linked to a target statistic.

The first parameter is \tilde{B} , which captures the disutility of work, as defined above. The most relevant target moment is the employment rate of 78.2% from the SIPP sample.¹⁴ The next parameter β captures the discount factor of households and is targeted to match a quarterly interest rate of 1%. The next parameter σ_x governs the variability of idiosyncratic labor productivity. We calibrate it to match the overall wage dispersion captured by the Gini index of worker wages. The target statistic is chosen to be 0.359, which is the average Gini wage in 2000 (Heathcote,

¹²This is the estimate for when the Earned Income Tax Credit (EITC) is included because we do not consider it in our calibration of welfare transfers. We also considered alternative values for λ_p , but these did not affect our quantitative results substantially. This quantitative insignificance of tax progressivity can also be seen explicitly in our counterfactual exercise in Section 5.3.

¹³An interesting exercise for the future would be to investigate how our results might carry over in the presence of other types of aggregate shocks—on top of the standard TFP shocks. The estimation of multiple aggregate shocks within a model including heterogeneous agents and nonconvexities is an important task yet is difficult at the present moment due to computational costs.

¹⁴This value is higher than the employment-population ratio of around 60% that was used in the previous literature (e.g., Chang and Kim 2007) because we focus here on working-age samples.

Table 1: Parameter values chosen internally

Values	Parameters		Target statistics		
		Description	Model	Data	Description
$\tilde{B} =$.692	Disutility of work	.777	.782	Employment rate
$\beta =$.985	Subject discount factor	.010	.010	Real interest rate
$\sigma_x =$.126	S.D. of innovations to $\ln x$.360	.359	Wage Gini index
$T_1 =$.0337	Overall transfer size	.044	.044	Ratio of Avg ($T_1 + T_2$) to output
$\omega_s =$.117	Scale of non-flat transfers	.0203	.0201	Ratio of Avg T_2 to output
$\omega_p =$	3.62	Progressivity of transfers	3.07	3.06	$E(T_2 1st\ income\ quintile)/E(T_2)$

Perri, and Violante 2010).¹⁵

The last three parameters— T_1 , ω_s and ω_p —govern statistics related to transfers. Recall that T_1 determines the size of universal transfers and ω_s determines the scale of non-flat transfers (T_2). Therefore, the first target statistic regarding transfers is set as the total transfers-output ratio of 4.4%. This is obtained from the time-series average of the ratio of transfers (excluding Social Security and Medicare) to output over the years 1961-2016 according to the Bureau of Economic Analysis (BEA) data. The next target is the average government expenditures on social benefits related to income-security (Table 3.12 from the BEA) over the years 1961–2016—that is, 2.0% of output.¹⁶ Next, ω_p shapes the degree of progressivity in government transfers. Our calibration strategy is to let the model replicate an empirically reasonable degree of transfer progressivity through ω_p , given the value of ω_s . For this purpose, we measure the degree of progressivity in the U.S. transfer programs using the SIPP data. We construct a broad measure of government transfers, including means-tested programs and social insurance (as detailed in Appendix E). Since these welfare programs are highly relevant for poor households, we choose as a target statistic the ratio of the average amount of means-tested transfers received by the first income quintile to its unconditional mean (3.06) (Yum 2018).

Calibrating alternative model specifications Having explained the calibration strategy of our baseline model, Model (HA-T), we now describe how we calibrate our nested model specifications: Model (HA-N), Model (HA-F), and Model (RA). In general, it would be ideal to minimize the number of parameters to be recalibrated to ensure that our comparison across

¹⁵Inequality has been steadily rising in the U.S. In Appendix H, we also consider different values for this target.

¹⁶We select the components in order to be consistent with our measurement of transfers from the SIPP data, as described below. Our classification of transfers is similar to Krusell and Rios-Rull (1999). See Appendix E for details.

different model specifications is not driven by different values of parameters. We therefore hold as constant across the nested specifications the parameters governing the idiosyncratic productivity risk: ρ_x and σ_x . However, it is necessary to recalibrate a subset of the internally calibrated parameters to ensure that the different model specifications are similar in terms of their target statistics (e.g., the employment and interest rates in steady state equilibrium).

First, consider Model (HA-N). Because it abstracts from transfers ($T_1 = \omega_s = 0$), the parameter ω_p is irrelevant. Thus, we recalibrate \tilde{B} and β to match the employment rate of 78.2% and the real interest rate of 1%. This leads to values of $\tilde{B} = 0.974$ and $\beta = 0.9833$. Next, consider Model (HA-F), which shuts down differences in transfers across households. With $\omega_p = 0$, distinguishing between T_1 and ω_s becomes unnecessary. Therefore, we can calibrate the sum of T_1 and ω_s to match the total transfers-output ratio of 4.4%. Aside from this, we recalibrate \tilde{B} and β in the same manner, to yield values of 0.714 and 0.9848, respectively. Finally, and unlike the heterogeneous-agent models, Model (RA) can be calibrated analytically, as is shown in Appendix B. As for the parameters related to tax and transfers, we simply use the average tax rate and the average transfers because progressivity is irrelevant in Model (RA).

Steady state properties Table 1 reports that the baseline model does a good job of matching the target statistics, and the other nested model specifications do a great job of matching a smaller number of targets as well. This does not necessarily mean that the model can account for other relevant statistics. We therefore present the (non-targeted) distributional aspects in steady state. First, Table 2 summarizes the share of wealth and the employment rates by wealth quintile from both the model and the data.¹⁷ Overall, all heterogeneous-agent model specifications do a good job of accounting for the shares of wealth held by each wealth quintile.

When we look at the employment rate by wealth quintile reported (also reported in Table 2), we can clearly see that Model (HA-T) does a significantly better job of accounting for the cross-sectional employment-wealth relationship. In the U.S., the employment rate of the first wealth quintile is relatively low (70.0%) and is then relatively flat across the other wealth quintiles. This weak inverted-U-shape of the employment rates across wealth quintiles in the data is relatively well captured in Model (HA-T). On the other hand, Model (HA-N) predicts that employment falls sharply with wealth, consistent with the findings of Chang and Kim (2007). In this class of the incomplete-markets framework, the existence of transfers mitigates the excessively strong precautionary motive for employment among poor households who expect to be near the borrowing limit in the near future (Yum 2018). Although Model (HA-F) mitigates the

¹⁷Table 2 presents statistics on wealth distribution obtained from the 1992–2007 Survey of Consumer Finances (SCF), as reported by Yum (2018). These statistics from the SCF show a greater concentration of wealth in the top wealth quintile because it better captures the top of the distribution by over-sampling the rich.

Table 2: Characteristics of wealth distribution

	Wealth quintile				
	1st	2nd	3rd	4th	5th
<i>Share of wealth (%)</i>					
U.S. Data (SIPP)	-2.2	1.2	6.8	18.4	76.3
U.S. Data (SCF)	-0.4	1.2	5.1	13.6	80.5
Model (HA-T)	-0.0	0.9	5.2	19.7	74.3
Model (HA-N)	-0.1	0.1	4.8	20.4	74.8
Model (HA-F)	-0.0	0.3	4.9	20.2	74.7
<i>Employment rate (%)</i>					
U.S. Data (SIPP)	70.0	77.9	80.9	82.5	79.7
Model (HA-T)	85.3	79.3	84.4	75.2	64.2
Model (HA-N)	100.0	99.2	74.0	66.0	51.9
Model (HA-F)	100.0	92.0	75.2	67.9	54.0

Note: U.S. data are based on the 2001 Survey of Income and Program Participation (SIPP) and the 1992–2007 Survey of Consumer Finances (SCF) (Yum 2018). Model (HA-T) is the heterogeneous-agent model with targeted transfers. Model (HA-N) is the heterogeneous-agent model with no transfers. Model (HA-F) is the heterogeneous-agent model with flat transfers.

negative wealth gradients in employment seen in Model (HA-N), it does not generate the nearly flat employment rates across wealth quintiles.

Next, Table 3 shows the micro relationship between income or employment status and transfers in the steady state equilibrium. Specifically, the reported numbers are the ratios of average progressive-component transfers in each income quintile (left panel) or in each employment status (right panel) to the unconditional mean progressive-component transfers. In the U.S., there is a clear negative relationship between income and the amount of income-security transfers. Note that, in our model, this is a complicated equilibrium object: it is shaped not only by the parametric assumption on the nonlinear transfer schedule (9) but also by endogenous household choices (such as those regarding consumption-saving and labor supply). Despite the relatively simple functional form (9), we can see that our baseline model does an excellent job of replicating the degree of transfer progressivity in the U.S. Note that, since differences in transfers across households are removed by design in Model (HA-F), this ratio is one for all income quintiles. Another way to examine the speed at which transfers phase out is to compute average welfare transfers by employment status. Although we do not target this directly, our baseline model generates a substantially higher ratio for non-workers (2.69) relative to workers (0.51), in line with the data.

Table 3: Progressivity of income-security transfers

	Income quintile					Full-time emp.	
	1st	2nd	3rd	4th	5th	Yes	No
<i>Conditional mean/unconditional mean</i>							
U.S. Data	3.06	0.99	0.52	0.26	0.17	0.48	2.86
Model (HA-T)	3.07	1.07	0.56	0.24	0.06	0.51	2.69
Model (HA-F)	1.00	1.00	1.00	1.00	1.00	1.00	1.00

Note: The source of U.S. data is the Survey of Income and Program Participation 2001.

We now present results related to the aggregate employment responses in steady state. In this regard, it is useful to consider the recent empirical evidence of Mui and Schoefer (2020), who propose a novel concept called the reservation raise. Specifically, a reservation raise, ξ , is defined as the ratio of the reservation wage, which would make an agent indifferent between the two choices of working and not working, to the actual (or potential) wage.¹⁸ Once computing these micro reservation raises, we then obtain their cumulative distribution function, denoted by $\mathcal{R}(\xi)$. As detailed in Mui and Schoefer (2020), if we consider a change in the aggregate raise Ξ , the arc elasticities of extensive-margin labor supply can be defined as

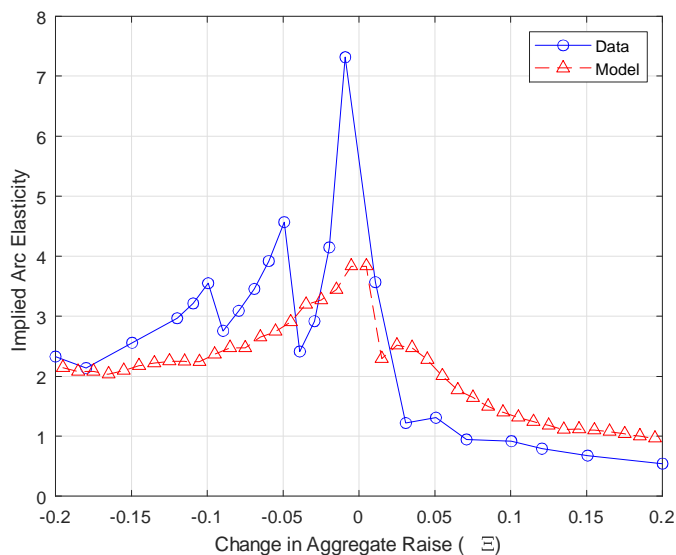
$$\varepsilon(\Xi) = \left(\frac{\mathcal{R}(1 + \Xi) - \mathcal{R}(1)}{\mathcal{R}(1)} \right) / \Xi,$$

with $\mathcal{R}(\xi)$ being interpreted as an aggregate labor supply curve along the extensive margin. Figure 2 shows $\varepsilon(\Xi)$ from Mui and Schoefer (2020) estimated using the representative U.S. sample as well as its model counterpart.

We find that our baseline model replicates two salient patterns observed in the data. As noted by Mui and Schoefer (2020), the empirical arc elasticities show that (i) local elasticities are large with respect to small changes in potential earnings (or reservation raises), and that (ii) elasticities are smaller with respect to large changes in potential earnings. It is worth noting that our baseline model not only qualitatively reproduces these two salient patterns, but it also generates empirically reasonable quantitative responses. Specifically, the elasticities are as high as three (or above) for very small changes in reservation raises (such as $\pm 5\%$). Moreover, they become smaller (at around two) when the change in potential earnings approaches -20% , and they get even smaller again (at around one) when the change in raises approaches $+20\%$. To sum up, it is reassuring to see that our model generates an empirically reasonable aggregate labor

¹⁸Hence, reservation raises for those who choose to work would be less than one (or $\xi < 1$), whereas those for non-workers would be larger than one (or $\xi > 1$).

Figure 2: Arc elasticities: data versus model



Note: The data plot is from Mui and Schoefer (2020) who report the U.S. estimates of arc elasticities. The arc elasticities are computed, based on the reservation raise distribution.

supply curve, which is a complicated, non-targeted object.¹⁹

5 Quantitative analysis

In this section, we report the main business cycle results and illustrate the mechanism underlying our main quantitative results.

5.1 Business cycle properties

We first compare the business cycle statistics of key macroeconomic variables from model simulations to those from the data. We filter all the series using the Hodrick–Prescott filter with a smoothing parameter of 1,600. The U.S. data statistics are computed using the aggregate data from 1961Q1 to 2016Q4 (see Appendix D for more details). Table 4 summarizes the cyclical volatility of the following key aggregate variables: output Y , consumption C , investment I , aggregate efficiency unit of labor L , aggregate hours H , and average labor productivity Y/H . Volatility is measured using the percentage standard deviation. As is standard in the business

¹⁹In Appendix H, we report aggregate labor supply curves and arc elasticities from the other heterogeneous-agent model specifications.

Table 4: Volatility of aggregate variables

	U.S. data	Model			
		(HA-T)	(HA-N)	(HA-F)	(RA)
σ_Y	1.50	1.27	1.48	1.46	1.83
σ_C/σ_Y	0.58	0.27	0.28	0.27	0.25
σ_I/σ_Y	2.96	2.87	2.99	2.99	3.08
σ_L/σ_Y	-	0.50	0.64	0.62	-
σ_H/σ_Y	0.98	0.73	0.51	0.60	0.80
$\sigma_{Y/H}/\sigma_Y$	0.52	0.64	0.54	0.57	0.25

Note: See Table 2 or Section 3.2 for the description of the model specifications. Each quarterly variable is logged and detrended using the Hodrick-Prescott filter with a smoothing parameter of 1600. Volatility is measured by the percentage standard deviation of each variable. The U.S. statistics are based on aggregate time-series from 1961Q1 to 2016Q4.

cycle literature, our discussion focuses on relative volatility, which is computed as the absolute volatility of each variable divided by that of output.

The most notable finding in Table 4 is that the high volatility of aggregate hours relative to output observed in U.S. data ($\sigma_H/\sigma_Y = 0.98$) is well accounted for by Model (HA-T). Standard real business cycle models are known to have difficulties in generating a large relative volatility of hours without relying on a low curvature of the utility function (or a high Frisch elasticity). When the stand-in household’s utility function in Model (RA) features zero curvature in labor supply, it indeed generates a substantial relative volatility of hours (0.80). It is striking that our baseline model, Model (HA-T), delivers a comparably high volatility of hours (0.73).

The results of Chang and Kim (2006, 2007) suggest that a large relative volatility of hours obtained through indivisible labor (Rogerson 1988) in Hansen (1985) may not be robust in incomplete markets economies with heterogeneous households. We can also see this point when we look at the performance of our Model (HA-N), which delivers a substantially smaller volatility of hours (0.51). This is the case even though our Model (HA-N) has progressive taxes—an additional feature relative to Chang and Kim (2006, 2007). However, as noted above, once heterogeneity in transfers is incorporated in line with the observed patterns in the micro data, our baseline model shows that the heterogeneous-agent incomplete markets model can perform similarly to the Hansen–Rogerson economy—in terms of a large volatility of hours over the business cycle.

The performance of Model (HA-F) reveals that introducing flat transfers into the model can help obtain a larger relative volatility of hours (0.60). However, this value is still quite far from its counterpart in Model (HA-T) of 0.73, suggesting that the phasing-out of transfers plays an

Table 5: Cyclicalities of aggregate variables

	U.S. data	Model			
		(HA-T)	(HA-N)	(HA-F)	(RA)
$Cor(Y, C)$	0.81	0.85	0.85	0.84	0.84
$Cor(Y, I)$	0.90	0.99	0.99	0.99	0.99
$Cor(Y, L)$	-	0.92	0.96	0.96	-
$Cor(Y, H)$	0.86	0.77	0.95	0.87	0.99
$Cor(Y, Y/H)$	0.30	0.69	0.95	0.85	0.84
$Cor(H, Y/H)$	-0.23	0.07	0.81	0.48	0.74

Note: See Table 2 or Section 3.2 for the description of the model specifications. Each quarterly variable is logged and detrended using the Hodrick-Prescott filter with a smoothing parameter of 1600. Cyclicalities are measured by the correlation of each variable with output. The statistics are based on aggregate time-series from 1961Q1 to 2016Q4.

important role. This finding is in line with our analytical results in Proposition 3 of Section 2, which showed that greater progressivity increases the degree to which aggregate hours fluctuates with respect to an aggregate shifter (TFP in this case).²⁰

We now move on to the cyclicalities of macroeconomic variables—a key focus of this paper. The first five rows of Table 5 show correlations between output and other aggregate variables. The last row shows the correlation between aggregate hours and labor productivity. As is well known in the literature (King and Rebelo 1999), most macroeconomic variables like consumption, investment, and aggregate hours are highly procyclical in the U.S. Table 5 shows that the strongly positive correlations with output are fairly well replicated in all of our model specifications, regardless of the presence of heterogeneity or institutional details. Therefore, one might conclude that heterogeneity or government transfers are irrelevant, at least with respect to the cyclicalities of macroeconomic variables over the business cycle.

However, we can see that this conclusion is premature when we look at the comovement of average labor productivity and output. In the U.S., strong procyclicalities are not a feature of average labor productivity (i.e., $Cor(Y, Y/H) = 0.30$). A related observation is that the correlation between hours and average labor productivity is even weakly negative (-0.23), often referred to as the Dunlop–Tarshis observation (Christiano and Eichenbaum 1992). By contrast,

²⁰There are other interesting differences in the volatility of macroeconomic aggregates. For instance, the volatility of average labor productivity over the business cycle tends to be more consistent with the data in the heterogeneous-agent models, as compared to Model (RA). Another observation is that the presence of government transfers tends to reduce the volatility of consumption over the business cycle. This suggests that government transfers play a role as stabilizers, effectively providing insurance against aggregate risk (e.g., see McKay and Reis 2016).

canonical real business cycle models generate highly procyclical average labor productivity, and thus fail to replicate the limited cyclicity of average labor productivity seen in the data. The high correlation between output and average labor productivity in Model (RA) (0.84) is also a manifestation of this weakness.²¹

The most notable finding in Table 5 is that the strong procyclicality of average labor productivity is considerably muted (0.69) in Model (HA-T), and as such it is closer to the data (0.30). In contrast to the existing literature—which tends to rely on the introduction of additional exogenous shocks (e.g., Benhabib et al. 1991; Christiano and Eichenbaum 1992; Braun 1994; and Takahashi 2020)—the key to our result is the interaction between household heterogeneity and transfer progressivity. This in turn generates heterogeneous labor supply behavior across households, as highlighted in Section 2. The importance of the interplay between household heterogeneity and transfers can be seen by the performance of our nested model specifications. Once we abstract from either household heterogeneity or differences in transfers across households, our model generates highly procyclical average labor productivity (above 0.8). In particular, when we abstract from transfers in their entirety (as in Chang and Kim, 2006, 2007) with Model (HA-N), we generate a very high correlation of 0.95. This implies that heterogeneity per se does not dampen highly procyclical average labor productivity in real business cycle models.

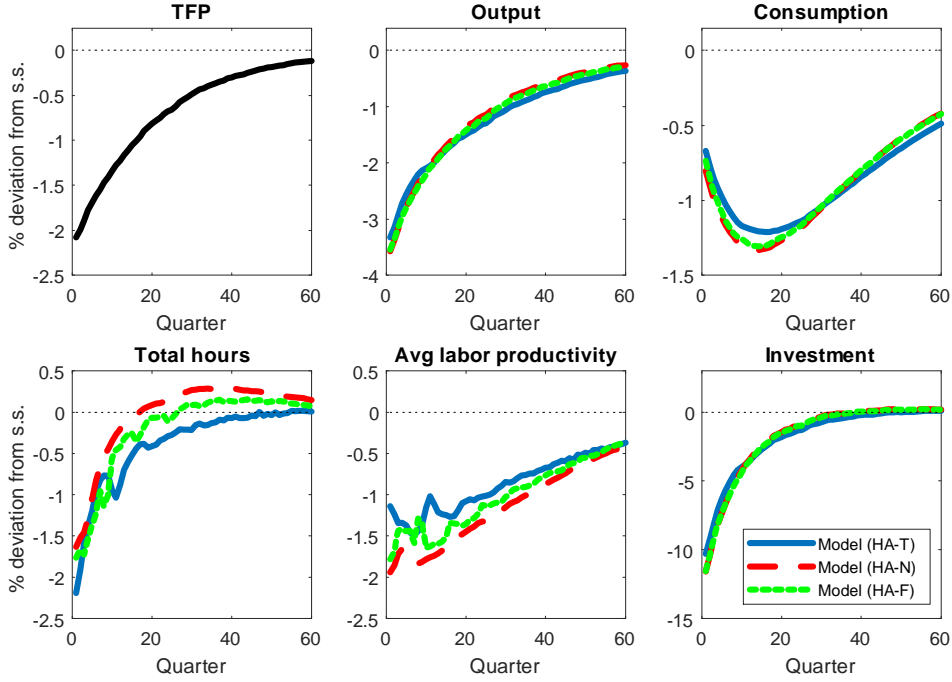
5.2 Impulse responses

We now investigate the mechanism underlying our quantitative success by using impulse response functions. Figure 3 shows the impulse responses of the key aggregate variables such as output, consumption, aggregate hours, average labor productivity, and investment following a persistent negative 2% shock to z (or TFP) for each of our heterogeneous-agent model specifications. We follow the simulation-based methodology developed by Koop et al. (1996), as described in detail in Appendix G (see also Bloom et al. 2018).

The impulse response of aggregate hours clearly confirms that Model (HA-T) (solid line) delivers a larger fall in hours than the nested heterogeneous-agent models—Model (HA-N) (dashed line) and Model (HA-F) (dotted line)—despite the fact that its output declines the least strongly on impact. Another important difference is the impulse responses of average labor productivity. In Model (HA-N), the dynamics of average labor productivity closely follow the pattern of output since it falls quite sharply on impact. This explains the very high correlation of Y/H with Y in Table 5. When flat transfers are present in Model (HA-F), we see that the overall decrease in average labor productivity is mitigated. In Model (HA-T), the magnitude of the fall in average

²¹These correlations would become even higher in models without indivisibility of labor (Hansen 1985) or in the absence of labor taxes.

Figure 3: Impulse responses of macroeconomic aggregates



Note: TFP denotes total factor productivity. The figures display the IRFs of macroeconomic aggregates to a negative 2 percent TFP shock with persistence ρ_z .

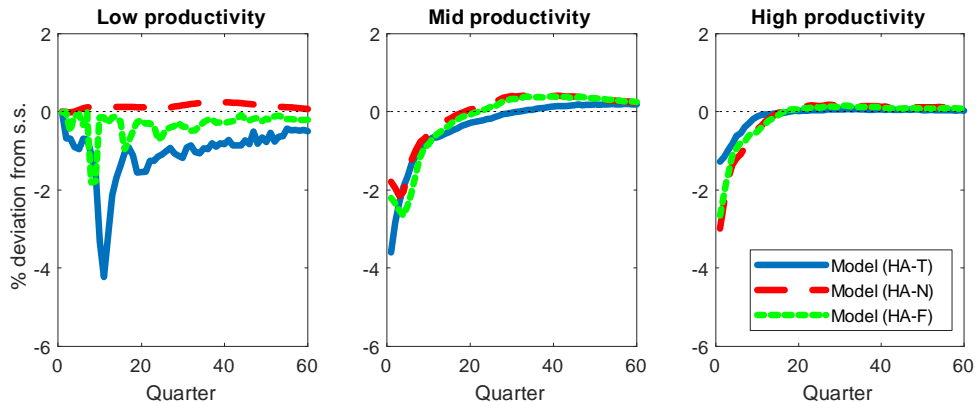
labor productivity is even smaller, despite it having the largest fall in hours.

To understand the underlying cause of these differences in aggregate dynamics, it is useful to investigate the impulse responses at a more disaggregated level.²² Specifically, in each period, we categorize households into three almost evenly distributed groups: (i) the low productivity group $\{x_i\}_{i=1}^4$; (ii) the mid productivity group $\{x_i\}_{i=5}^6$; and (iii) the high productivity group $\{x_i\}_{i=7}^{10}$. Figure 4 plots the impulse responses of hours by productivity following the same negative shocks, whereas Figure 5 plots its counterparts with respect to positive TFP shocks.

There are several important patterns worth noting. First, there is a relatively small difference in labor supply responses among the mid productivity group across the different model specifications. On the other hand, the response of the high productivity group is clearly weaker in Model (HA-T) compared to the other heterogeneous-agent models. Second, recall that Proposition 1 from our simple model implies that agents with lower potential earnings tend to be more elastic in

²²Another obvious candidate is the dynamics of equilibrium prices. Figure A5 displays changes in the market-clearing wage per efficiency units of labor and in real interest rates following the same negative TFP shock for each of our model specifications. It appears as though the difference between these specifications is not substantial among the heterogeneous agent models, suggesting that our main results are not driven mainly by the difference in equilibrium price dynamics.

Figure 4: Impulse responses of total hours by productivity



Note: Households are grouped into low productivity (below median), mid productivity (median), and high productivity (above median). The figures display impulse responses for employment in each group to a negative 2 percent TFP shock with persistence ρ_z .

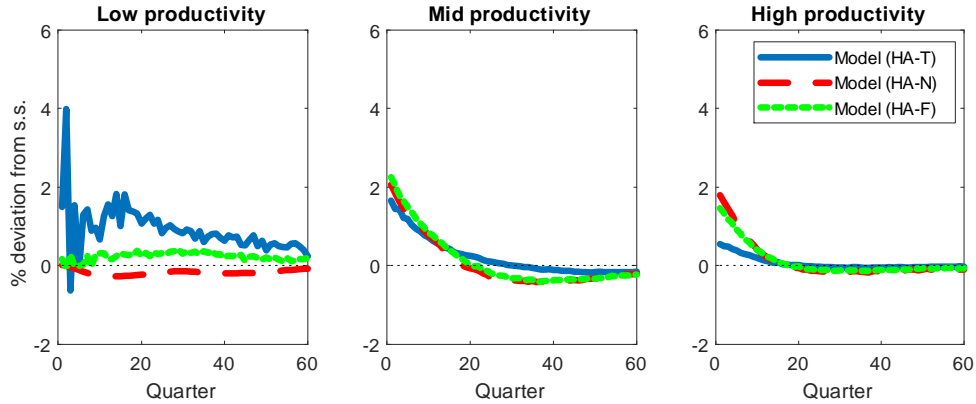
their labor supply. In fact, this pattern clearly applies to Model (HA-T), which generates greater magnitudes of changes in labor supply among low productivity groups. This heterogeneity of labor supply responses explains why Model (HA-T) is able to reduce the cyclical volatility of average labor productivity.

However, this monotonous relationship between elasticity and individual productivity breaks down for the low productivity group, especially in Model (HA-N). This exceptionally inelastic employment response is related to the results of Domeij and Floden (2006) and of Yum (2018). Specifically, both of these papers consider incomplete markets models without public insurance and show that wealth-poor households who lack self-insurance have precautionary labor supply motives at the intensive margin (Domeij and Floden 2006) and at the extensive margin (Yum 2018). Such precautionary motives can dominate the standard intertemporal substitution motive, which in turn could weaken the responses of hours with respect to a persistent fall in wages.

This inelastic labor supply among the low productivity group provides a key reason for both the lower volatility of total hours and the highly procyclical average labor productivity seen in Model (HA-N). This illustrates why heterogeneity per se is not sufficient to explain our key results in incomplete markets environments. The impulse responses from Model (HA-F)—which provides uniform transfers—show that the low productivity group has now become responsive to aggregate TFP changes though its magnitude is weak.

The precautionary motive would be less relevant when positive aggregate shocks hit the economy since they would move agents away from the borrowing limit. Figure 5 shows that, with respect to positive aggregate shocks, low productivity agents in Model (HA-T) do not show

Figure 5: Impulse responses of total hours by productivity with respect to positive TFP shocks



Note: Households are grouped into low productivity (below median), mid productivity (median), and high productivity (above median). The figures display impulse responses for employment in each group to a positive 2 percent TFP shock with persistence ρ_z .

mutated responses on impact. This is in contrast to their sluggish immediate responses following negative aggregate shocks in Figure 4. Since negative shocks tend to move agents toward the borrowing limit, they raise precautionary motives and thus weaken intertemporal substitution motives.²³

5.3 Progressivity and aggregate fluctuations

We now use our baseline model to conduct counterfactual exercises about how higher progressivity affects aggregate fluctuations. The first counterfactual adjusts the parameters in the transfer function so that higher overall tax-and-transfer progressivity is achieved by affecting low-income households disproportionately. We also consider a case where we raise the progressivity in the tax function while keeping the transfer function unchanged. To control for the strength of each policy reform, we ensure that each policy increases the difference between the income Gini coefficients before and after taxes and transfers by 2 percentage points, as compared to the baseline model.²⁴

Table 6 reveals that the first counterfactual exercise indeed increases the amount of transfers going to low-income households with a higher rate at which transfers are phased out. Conse-

²³Figure 5 also reveals that the overall magnitude of employment responses features asymmetry. Specifically, the overall responses are stronger with respect to negative shocks, as compared to those with respect to positive shocks. This is in fact in line with the distribution of arc-elasticities, based on reservation raises: Figure 2 shows that downward adjustments induce greater elasticities (above two) relative to upward adjustments (around one).

²⁴Specifically, tax progressivity is increased by raising λ_p by 80% (or $\lambda_p = 0.0954$). A higher λ_p tends to raise overall tax revenues, and this works to increase redistribution. As for transfer progressivity, we adjust both ω_p and ω_s simultaneously. This is because a higher ω_p tends to reduce the overall size of transfers, which works against an increase in redistribution. The required percentage increase is 25%.

Table 6: Effects of progressivity on the steady-state economy and aggregate fluctuations

	Baseline Model (HA-T)	Higher progressivity	
		Transfer function	Tax function
Steady state			
<i>Employment rate (%)</i>			
Overall	77.7	71.2	78.4
<i>By wealth quintile</i>			
1st	85.3	50.0	92.7
2nd	79.3	83.8	75.3
3rd	84.4	80.6	85.1
4th	75.2	76.8	74.6
5th	64.2	64.6	64.2
 <i>- Cond. mean/uncond. mean of T_2 by income quintile</i>			
1st	3.07	3.61	3.03
2nd	1.07	0.92	1.09
3rd	0.56	0.34	0.57
4th	0.24	0.12	0.25
5th	0.06	0.02	0.06
 Business cycles			
σ_Y	1.27	1.37	1.29
σ_H/σ_Y	0.73	1.09	0.75
$Cor(Y, Y/H)$	0.69	0.19	0.66
$Cor(H, Y/H)$	0.07	-0.44	0.05

Note: Each counterfactual exercise leads to the same Gini of after-tax-and-transfer income using each policy instrument.

quently, the steady-state employment rates of wealth-poor households are affected heavily. On the other hand, the second counterfactual exercise has more balanced effects on the employment rates across the distribution, and the pattern of transfers across the income distribution remain nearly unchanged.

Having checked the steady state effects, we now move on to business cycle implications. Table 6 reports that the first counterfactual exercise reduces the cyclicity of average labor productivity and raises the volatility of total hours quite substantially. On the other hand, the second counterfactual exercise has limited effects on aggregate labor market dynamics. Average labor productivity becomes slightly less procyclical and the volatility of hours increases marginally.²⁵

²⁵In Figure A3, we plot aggregate labor supply curves (or reservation raise cumulative distributions) and their

Overall, the above results reinforce our conclusion in the previous subsection that changes in the rate at which transfers are phased out are quantitatively more important than changes in average tax progressivity when it comes to aggregate labor market fluctuations with labor supply at the extensive margin.

6 Microeconomic evidence of heterogeneity in the extensive margin labor supply responses

As shown in the previous sections, the key mechanism of our model relies on heterogeneous labor supply responses. More precisely, households with low potential earnings are considerably more elastic in adjusting their labor supply at the extensive margin, which weakens a highly procyclical average labor productivity and enlarges the volatility of aggregate hours worked over the business cycle. In this section, we empirically document heterogeneity in labor supply responses to verify whether our key model mechanism exists in the micro data.²⁶

Specifically, we exploit the panel structure of the PSID to explore whether extensive margin labor supply responses differ as a function of hourly wage. This panel structure is useful because we can keep track of the same people and observe their labor supply decisions over time. Because labor supply changes can be measured in different ways and can be shaped by forces at different levels (i.e., idiosyncratic vs. aggregate), we consider two approaches. The first approach computes the probability of the extensive margin labor supply adjustment for each individual and illustrates how it differs by wage. The second approach focuses on differences in the magnitude of full-time employment rate changes across wage groups during the last six recessions.

The first approach requires us to have relatively long time-series observations for each individual to obtain a consistent estimate of the adjustment probability, based on individual-level flow data.²⁷ First, let us fix the year at j , and denote i as an individual index, and t the year when the individual is observed. We define the extensive margin adjustment based on a full-time employment indicator $E_{i,t}$: an individual i in year t is in full-time employment (i.e., $E_{i,t} = 1$) if the annual number of hours worked is greater than 1,000.²⁸ Then, we define a binary switching

corresponding arc elasticities for the two counterfactual experiments, which corroborate our cyclical volatility results.

²⁶There is limited empirical evidence of heterogeneity in labor supply responses at the extensive margin across wage groups. See Kydland (1984) and Juhn et al. (1991) for earlier evidence. Hoynes et al. (2012) provide evidence of heterogeneity in employment rates across other (potentially related) dimensions, such as race, gender, age, and education.

²⁷Since the frequency of the PSID survey had been annual until 1997 and became biannual from 1999 onward, we use only samples observed annually from the 1969–1997 waves.

²⁸The results in this section are quite robust to alternative threshold values for the full-time employment variable. In Appendix I, we report the results when we use 1,500 hours as a full-time threshold value.

Table 7: Probability of extensive margin adjustment, by wage quintile

	The length of tracking time T								
	5 years			10 years			15 years		
<i>Wage quintile</i>	Switches			Switches			Switches		
<i>in base year</i>	All	Pos only	Neg only	All	Pos only	Neg only	All	Pos only	Neg only
1st	.097	.061	.036	.075	.048	.027	.066	.042	.024
2nd	.051	.030	.020	.042	.025	.017	.038	.022	.015
3rd	.038	.020	.018	.032	.018	.014	.031	.017	.013
4th	.034	.016	.018	.028	.014	.015	.026	.012	.014
5th	.037	.018	.019	.032	.015	.017	.030	.014	.016
Base years	1969–1993 ($J = 25$)			1969–1988 ($J = 20$)			1969–1983 ($J = 15$)		
Avg. no. obs in base years	1,677			1,189			834		
Total no. obs.	41,920			23,783			12,514		
Avg. age	40.2			41.0			41.5		

Note: See text for the definition of the switching probability reported in this table. Numbers in parentheses show the number of base years. We use samples whose age is between 22 and 64 (inclusive) and who are heads and are not self-employed. "All" refers to the baseline estimates when using both positive and negative switches, whereas "pos only" and "neg only" use only positive ones (i.e., $E_{i,t} = 1$ and $E_{i,t-1} = 0$) and only negative ones (i.e., $E_{i,t} = 0$ and $E_{i,t-1} = 1$), respectively.

variable $S_{i,t}$ such that it equals one if $E_{i,t} \neq E_{i,t-1}$ and it equals zero otherwise. We exclude transitions from $E_{i,t-1} = 1$ to $E_{i,t} = 0$ if the individual has a non-zero unemployment spell in period t in order to rule out transitions caused by layoffs.

Note that, given the length of time spent tracking each individual T , there are $T - 1$ counts of $S_{i,t}$ for each individual i . Once we take the average over time, we obtain the individual-specific probability of an extensive-margin adjustment with an annual frequency (i.e., $p_{i,j} \equiv \frac{1}{T-1} \sum_{t=j+1}^{j+T-1} S_{i,t}$). As we are interested in differences across the wage distribution, we compute p_j^q , which is defined as the conditional mean of $p_{i,j}$ for the wage quintile bin of each individual $q \in \{1, 2, \dots, 5\}$ determined in the base year j .

Because different values for the length of time spent tracking each individual entails a trade-off, we consider three variants: $T \in \{5, 10, 15\}$. On the one hand, a larger number is beneficial because we are more likely to have a consistent estimate of the adjustment probability at the individual level. On the other hand, a longer tracking time implies a tighter restriction on the samples—because we keep only samples that were observed for T consecutive years. Given

the value of T , we compute the estimates of $\{p_j^q\}_{q=1}^5$ by changing the base year j . That way, we attempt to mitigate variations due to differences in the initial wage distribution, which is potentially affected by business cycle fluctuations. The reported values in Table 7 are the mean switching probabilities for each wage quintile averaged across the base years, $p^q \equiv \frac{1}{J} \sum p_j^q$, where the number of base years J is reported in parentheses.

Table 7 reveals a clear pattern: the individual-level probability of adjusting the extensive margin is significantly higher among low-wage workers. For instance, when $T = 5$, the probability of switching to or from full-time employment among the first wage quintile is 9.7% (the annual frequency). In particular, we can see that this probability tends to decrease with wage. For the third to fifth quintiles, this probability is relatively flat at approximately 3.5%. When T increases, we also find that the key pattern of extensive margin adjustment probabilities across wage quintiles is still present. However, because the samples become slightly older and T becomes longer, we also see that their switching probabilities generally become lower.

We also compute these statistics using only either positive switches ($E_{i,t} - E_{i,t-1} > 0$) or negative switches ($E_{i,t} - E_{i,t-1} < 0$). Table 7 shows that negative wage gradients in the probability of full-time employment adjustments are present in both cases. Interestingly, positive switches feature not only higher adjustment probabilities but also a quantitatively larger negative wage gradient, showing that the overall gradient is more strongly driven by such positive adjustments.

The above exercise is based on long-run information regarding labor market flows at the individual level. The next empirical exercise instead uses the differences in magnitude of full-time employment level changes across wage groups during recessions. Specifically, we choose six recessions and for each recession we choose a peak year and a trough year, as guided by the cyclical component of quarterly real GDP per capita (Figure A1). Our definition of peak and trough years is limited by the frequency of the PSID because the data set was available annually until 1997 and only biannually since 1999. Therefore, our choice is also based on declines in aggregate employment during each recession event—according to our micro samples from the PSID. The resulting year combinations for each recession are shown in Table 8.

Next, we compute the conditional mean of full-time employment by wage quintile in the peak year for each recession by $\frac{1}{N_{peak}^q} \sum_i E_{i,peak}^q$, where N_{peak}^q is the number of observations in wage quintile bin q during the peak year. We then measure the percentage changes in the full-time employment rate by wage quintile in the corresponding trough year. It is important to note that we keep the set of households in each wage group fixed by assigning a wage quintile to each household in the peak year. That way, our measured changes by wage quintile are not affected by compositional changes, but are rather based on changes from within the same households.

Table 8 clearly shows that the employment rate fell most sharply in the first and second wage quintiles during all of the recessions, and that the magnitude of these declines tends to be

Table 8: Full-time employment changes in recessions, by wage quintile

	Recession					
	1969–71	1973–76	1980–83	1990–92	2000–02	2006–10
<i>Wage quintile in peak year</i>						
1st	-7.1	-10.1	-9.8	-8.7	-10.0	-17.2
2nd	-3.2	-7.4	-4.5	-6.3	-5.6	-12.9
3rd	-3.7	-7.0	-5.2	-4.6	-3.0	-11.2
4th	-4.7	-4.5	-5.9	-5.4	-4.8	-10.4
5th	-0.9	-5.7	-4.9	-4.6	-2.2	-5.8
No. obs.	1,655	1,756	2,007	2,166	2,924	2,802

Note: The full-time employment threshold is set to 1,000 annual hours. The year ranges denote the peak and trough years of each recession. Reported values are percentage changes in the full-time employment rate by wage quintiles (in the peak year of each recession) following the same set of individuals.

smaller among the higher wage quintiles. For example, the full-time employment rate among the first wage quintile during the last recession (i.e., the Great Recession) fell by 17.2%, whereas the counterpart among the fifth wage quintile fell by only 5.8%. This pattern of full-time employment changes by wage quintiles is quite robust across different recessions despite variations in overall magnitude.²⁹

One may be concerned about the possibility that the wage gradient of full-time employment changes found in Table 8 is driven largely by the demand channel of the firms, and that this may affect household employment status differentially across the wage distribution. To alleviate this concern, we also use the information from the PSID data about unemployment spells (available since the 1976 wave or the year of 1975) and exclude samples that experienced any unemployment spells in either the peak year or the trough year. We thereby attempt to rule out the effects caused by differential layoff probabilities across the wage distribution, although the number of observations in each recession decreases because of this additional sample restriction. Table A5 shows that although the magnitudes of full-time employment changes are somewhat weaker, the magnitude is negatively related to wage quintiles during most recessions.

Although the above two approaches are designed to capture different aspects of labor supply adjustments, they yield consistent results: the employment adjustments of lower-wage workers are more volatile. Both of these empirical findings are therefore consistent with the pattern of

²⁹Note that the overall magnitude of the fall in employment is relatively greater in the recessions of 1973–76, 1980–83 and 2006–10. This finding is, in fact, consistent with the relatively larger amplitudes of these recessions, as shown in Figure A1. This provides some external validation for our micro samples.

heterogeneity in labor supply responses seen in the baseline model. Nevertheless, we would also like to stress that the results in this section are only suggestive since we cannot rule out other possible demand-related effects from being behind the observed heterogeneous patterns. For example, they could be consistent with intrasectoral changes in the quality of workers made by firms (Ohanian 2001). Alternatively, there might be disproportionately more low-wage workers in those industries that are more vulnerable to the effects of the business cycle (Hoynes et al. 2012), or labor hoarding might be differential across wage groups.

7 Conclusion

In this paper, we have explored the interplay of household heterogeneity and tax-and-transfer progressivity in shaping the dynamics of macroeconomic aggregates over the business cycle. Using analytical results obtained from a stylized static model of the extensive margin labor supply, we first presented the key insight that higher progressivity would strengthen negative wage gradients in employment responses to aggregate shifters. Using a general equilibrium business cycle model with household heterogeneity, we have shown that micro-level heterogeneity substantially shapes the dynamics of aggregate labor market variables when heterogeneity interacts with progressive transfers. In particular, our baseline model delivers less procyclical average labor productivity when compared to the nested models that are similar to standard real business cycle models. At the same time, it retains the success of the canonical representative-agent indivisible labor supply model in generating a large volatility of aggregate hours without the assumptions of lotteries and perfect consumption insurance (Rogerson 1988). Our counterfactual analysis shows that in the U.S. tax and transfer system, the rate at which transfers are phased out is quantitatively very important for understanding labor supply along the extensive margin.

There are several future research questions that follow naturally from our study. One interesting and novel result that we highlighted in this paper is that the effects of higher progressivity can be quite different depending on which part of the income distribution is more affected by the fiscal change. A straightforward application would be to design a tax-and-transfer system that takes into account the welfare costs of business cycles in the presence of heterogeneous agents (e.g., Krusell et al. 2009). Second, although the current paper focuses on the total factor productivity shock, it would be interesting to explore how other types of aggregate shocks (such as monetary policy shocks) would be transmitted differently in our framework. Finally, our paper introduces some theoretical and quantitative mechanisms suggesting that changes in progressivity at the bottom of the income distribution (e.g., welfare transfers) might be behind the vanishing procyclicality of average labor productivity, as documented by Galí and van Rens

(2021).³⁰ Formal investigations of these changing relationship are out of the scope of the current paper, but they would nonetheless be highly valuable to address in future work.

References

- Ahn, SeHyoun, Greg Kaplan, Benjamin Moll, Thomas Winberry, and Christian Wolf. 2017. "When Inequality Matters for Macro and Macro Matters for Inequality." In *NBER Macroeconomics Annual 2017*, Volume 32: University of Chicago Press.
- Aiyagari, S. Rao. 1994. "Uninsured Idiosyncratic Risk and Aggregate Saving." *The Quarterly Journal of Economics* 109 (3): 659-684.
- 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.
- Benhabib, Jess, Richard Rogerson, and Randall Wright. 1991. "Homework in Macroeconomics: Household Production and Aggregate Fluctuations." *Journal of Political Economy* 6: 1166-1187.
- Ben-Shalom, Yonatan, Robert A. Moffitt, and John Karl Scholz. 2011. "An Assessment of the Effectiveness of Anti-Poverty Programs in the United States." NBER Working Papers No. 17042.
- Bick, Alexander and Nicola Fuchs-Schündeln. 2017. "Taxation and Labour Supply of Married Couples Across Countries: A Macroeconomic Analysis." *The Review of Economic Studies* 85 (3): 1543-1576.
- Bils, Mark J. 1985. "Real Wages Over the Business Cycle: Evidence from Panel Data." *Journal of Political Economy* 93 (4): 666-689.
- Bloom, Nicholas, Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten, and Stephen J. Terry. 2018. "Really Uncertain Business Cycles." *Econometrica* 86 (3): 1031-1065.
- Braun, R. Anton. 1994. "Tax Disturbances and Real Economic Activity in the Postwar United States." *Journal of Monetary Economics* 33 (3): 441-462.
- Chang, Yongsung and Sun-Bin Kim. 2014. "Heterogeneity and Aggregation: Implications for Labor-Market Fluctuations: Reply." *The American Economic Review* 104 (4): 1461-1466.

³⁰There has been a steady increase in the size of welfare programs according to the BEA data. For example, the total government social benefits we considered in our quantitative model (cash transfers, and food, medical, and childcare support to low income households) have increased relative to GDP from 0.5% in the 1960s to 3.5% in the early 2010s.

- . 2007. "Heterogeneity and Aggregation: Implications for Labor-Market Fluctuations." *American Economic Review* 97 (5): 1939-1956.
- . 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, 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.
- Chang, Yongsung, Sun-Bin Kim, and Frank Schorfheide. 2013. "Labor-Market Heterogeneity, Aggregation, and Policy (in) Variance of DSGE Model Parameters." *Journal of the European Economic Association* 11: 193-220.
- Christiano, Lawrence J. and Martin Eichenbaum. 1992. "Current Real-Business-Cycle Theories and Aggregate Labor-Market Fluctuations." *American Economic Review* 82 (3): 430-450.
- Conesa, Juan Carlos, Sagiri Kitao, and Dirk Krueger. 2009. "Taxing Capital? Not a Bad Idea After all!" *American Economic Review* 99 (1): 25-48.
- Doepke, Matthias and Michele Tertilt. 2016. "Families in Macroeconomics." In *Handbook of Macroeconomics*. Vol. 2, 1789-1891: Elsevier.
- Domeij, David and Martin Floden. 2006. "The Labor-Supply Elasticity and Borrowing Constraints: Why Estimates are Biased." *Review of Economic Dynamics* 9 (2): 242-262.
- Ferrière, Axelle and Gaston Navarro. 2022. "The Heterogeneous Effects of Government Spending: It's all about Taxes." Unpublished Manuscript.
- Ferriere, Axelle, Philipp Grubener, Gaston Navarro, and Oliko Vardishvili. 2022. "On the optimal design of transfers and income-tax progressivity." Unpublished Manuscript.
- Fleck, Johannes, Jonathan Heathcote, Kjetil Storesletten, and Giovanni L. Violante. 2021. "Tax and Transfer Progressivity at the US State Level." Unpublished Manuscript.
- Galí, Jordi and Thijs Van Rens. 2021. "The Vanishing Procyclicality of Labour Productivity." *The Economic Journal* 131 (633): 302-326.
- Guner, Nezih, Remzi Kaygusuz, and Gustavo Ventura. 2020. "Child-Related Transfers, Household Labour Supply, and Welfare." *The Review of Economic Studies* 87 (5): 2290-2321.
- . 2014. "Income Taxation of US Households: Facts and Parametric Estimates." *Review of Economic Dynamics* 17 (4): 559-581.
- Hansen, Gary D. 1985. "Indivisible Labor and the Business Cycle." *Journal of Monetary Economics* 16 (3): 309-327.

- Hagedorn, Marcus and Iourii Manovskii. 2008. "The Cyclical Behavior of Equilibrium Unemployment and Vacancies Revisited." *American Economic Review* 98 (4): 1692-1706.
- Heathcote, Jonathan. 2005. "Fiscal Policy with Heterogeneous Agents and Incomplete Markets." *The Review of Economic Studies* 72 (1): 161-188.
- Heathcote, Jonathan, Fabrizio Perri, and Giovanni L. Violante. 2010. "Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States, 1967–2006." *Review of Economic Dynamics* 13 (1): 15-51.
- Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante. 2020. "How should Tax Progressivity Respond to Rising Income Inequality?" *Journal of the European Economic Association* 18 (6): 2715-2754.
- . 2014. "Optimal Tax Progressivity: An Analytical Framework." *The Quarterly Journal of Economics* 132 (4): 1693-1754.
- . 2010. "The Macroeconomic Implications of Rising Wage Inequality in the United States." *Journal of Political Economy* 118 (4): 681-722.
- Holter, Hans A., Dirk Krueger, and Serhiy Stepanchuk. 2019. "How do Tax Progressivity and Household Heterogeneity Affect Laffer Curves?" *Quantitative Economics* 10 (4): 1317-1356.
- Hoynes, Hilary, Douglas L. Miller, and Jessamyn Schaller. 2012. "Who Suffers during Recessions?" *The Journal of Economic Perspectives* 26 (3): 27-47.
- Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes. 1995. "Precautionary Saving and Social Insurance." *Journal of Political Economy*: 360-399.
- Huggett, Mark. 1993. "The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies." *Journal of Economic Dynamics and Control* 17 (5): 953-969.
- Juhn, Chinhui, Kevin M. Murphy, and Robert H. Topel. 1991. "Why has the Natural Rate of Unemployment Increased Over Time?" *Brookings Papers on Economic Activity* 1991 (2): 75-142.
- Khan, Aubhik and Julia K. Thomas. 2008. "Idiosyncratic Shocks and the Role of Nonconvexities in Plant and Aggregate Investment Dynamics." *Econometrica* 76 (2): 395-436.
- King, Robert G. and Sergio T. Rebelo. 1999. "Resuscitating Real Business Cycles." *Handbook of Macroeconomics* 1: 927-1007.
- Koop, Gary, M. Hashem Pesaran, and Simon M. Potter. 1996. "Impulse Response Analysis in Nonlinear Multivariate Models." *Journal of Econometrics* 74 (1): 119-147.

- Krueger, Dirk, Kurt Mitman, and Fabrizio Perri. 2016. "Macroeconomics and Household Heterogeneity." *Handbook of Macroeconomics* 2: 843-921.
- Krusell, Per, Toshihiko Mukoyama, Ayşegül Şahin, and Anthony A. Smith Jr. 2009. "Revisiting the Welfare Effects of Eliminating Business Cycles." *Review of Economic Dynamics* 12 (3): 393-404.
- Krusell, Per and Jose-Victor Rios-Rull. 1999. "On the Size of US Government: Political Economy in the Neoclassical Growth Model." *American Economic Review*: 1156-1181.
- Krusell, Per and Jr Smith Anthony A. 1998. "Income and Wealth Heterogeneity in the Macroeconomy." *Journal of Political Economy* 106 (5): 867-896.
- Kydland, Finn E. 1984. "Labor-Force Heterogeneity and the Business Cycle." *Carnegie-Rochester Conference Series on Public Policy* 21 (1): 173-208.
- Kydland, Finn E. and Edward C. Prescott. 1982. "Time to Build and Aggregate Fluctuations." *Econometrica* 50 (6): 1345-1370.
- McKay, Alisdair and Ricardo Reis. 2016. "The Role of Automatic Stabilizers in the US Business Cycle." *Econometrica* 84 (1): 141-194.
- Mui, Preston and Benjamin Schoefer. 2020. "Reservation Raises: The Aggregate Labor Supply Curve at the Extensive Margin." Unpublished Manuscript.
- Ohanian, Lee E. 2001. "Why did Productivity Fall so Much during the Great Depression?" *American Economic Review* 91 (2): 34-38.
- Rogerson, Richard. 1988. "Indivisible Labor, Lotteries and Equilibrium." *Journal of Monetary Economics* 21 (1): 3-16.
- Takahashi, Shuhei. 2020. "Time-Varying Wage Risk, Incomplete Markets, and Business Cycles." *Review of Economic Dynamics* 37: 195-213.
- . 2014. "Heterogeneity and Aggregation: Implications for Labor-Market Fluctuations: Comment." *American Economic Review* 104 (4): 1446-1460.
- Yum, Minchul. 2018. "On the Distribution of Wealth and Employment." *Review of Economic Dynamics* 30: 86-105.