Employment Status and Credit Cycles: An Extensive Margin Analysis in Heterogeneous Agent Life Cycle Models.

Calvin Clack
Department of Economics
University of Illinois Chicago

August 19, 2026

Abstract

This paper addresses a critical gap between structural modeling conventions and empirical evidence regarding labor market adjustment margins. While most structural models emphasize the intensive margin (hours worked per agent), empirical evidence underscores the extensive margin (employment status), which remains overlooked even in models capable of capturing the dynamics it generates. This discrepancy is resolved through an analysis of a three period Overlapping Generations (OLG) model with exogenous youth unemployment and with a Heterogeneous Agent (HA) Life cycle framework, featuring age dependent unemployment distributions and precautionary savings motivations with exogenous labor productivity shocks.

The comparative statics of the OLG demonstrate that variations in the extensive margin fundamentally reshape credit markets by lowering household debt levels and real interest rates. Within the HA model, real household debt dropped by roughly 3.3% per 1% increase in unemployment and exhibited two key mechanisms: decreased loan demand from young unemployed agents facing borrowing constraints and increased supply from households self-insuring against elevated unemployment risk through precautionary savings. Impulse response analysis reveals that negative productivity shocks induce sharper household deleveraging when labor changes through the extensive margin. The theoretical implications are empirically validated using both Vector Autoregression and Local Projection methods. These findings extend the literature by establishing unemployment as an endogenous channel, particularly relevant to household deleveraging due to high unemployment, offering a foundation for integrating extensive margin labor dynamics into structural models.

Keywords: Life Cycle Models and Saving, Unemployment Rate, Precautionary Saving
JEL Classification: E21, E24, E44

1 Introduction

(Chetty, Guren, Manoli, and Weber 2011) found the aggregate hours macroeconomic Frisch elasticity to be 2.84, with the extensive margin Frisch elasticity being 2.30. Looking at US data during the Global Financial Crisis, the quarterly unemployment rate increased from 4.4% in the fourth quarter of 2006 to 9.9% in the first quarter of 2009. This corresponded to a deleveraging of real household debt, which decreased 14.1% from Q4 2007 to Q1 2013. Similar increases in unemployment and deleveraging followed across a majority of the globe. Cross-country variation is evident in Europe, with Greece, Italy, and Spain bearing the brunt of this GFC. Portugal saw its unemployment rate increase by 9.2% with real household debt deleveraging 27.4%, Italy saw its unemployment rate increase by 6.9% with real household debt deleveraging 8.3%, Greece saw its unemployment rate increase by 20.5% with real household debt deleveraging 30.9%, and Spain saw its unemployment rate increase by 18.3% with real household debt deleveraging 32.4%. The trend of rising unemployment and real household debt deleveraging motivates exploring unemployment as an endogenous channel, particularly relevant to household deleveraging.

Traditional macroeconomic models, such as the Real Business Cycle, New Keynesian, Heterogeneous Agent, and Heterogeneous Agent New Keynesian models, are not designed to capture the effects of inter-agent lending when labor fluctuates at the extensive margin. The labor market dynamic that they focus on is the intensive margin (adjusting an agent's hours worked) rather than the extensive margin (adjusting the number of agents working). The current literature on models with heterogeneous agents that vary over the extensive margin (employment) does not focus on household loans involving borrowing and lending between generations. I am studying the effect of changes along the extensive margins on household loan demand and loan supply in a continuous-time life-cycle model with incomplete markets. This enables the study of interest rates and loan dynamics in a manner that more closely reflects real-world behavior.

This paper makes three contributions to the literature on household deleveraging and labor market dynamics in structural models. First, empirical evidence identifies the extensive margin as an endogenous channel for household debt adjustment: both the unemployment rate and real household debt are statistically significant, whereas weekly hours worked is not. This is expanded by running different splits based on high vs. low levels of variation in weekly hours worked, unemployment, real household debt, and the household debt-to-GDP ratio to demonstrate how countries that exhibit different characteristics respond to shocks. Second, a 3-period OLG model identifies the employment based borrowing constraints, linking youth unemployment to decreased loan demand. Third, a more complex heterogeneous-agent life cycle model demonstrates why real household debt decreases when negative productivity shocks increase unemployment due to the borrowing constraint and despite an increased precautionary saving motive.

The analysis begins by examining a variance decomposition of total hours worked and the relationship between employment and real household Debt to establish the empirical motivation. European panel data will be used because it provides substantial cross-country variation not available in U.S. state-level observations. The data will be detrended using a method similar to (Hansen 1985), one of the first papers to perform labor market variance decomposition and model the RBC with labor fluctuations over the extensive margin, to decompose the variance of total hours worked in the economy. The analysis revealed that fluctuations were primarily driven by the extensive margin (employment). The intensive margin (weekly hours worked per worker) played a much smaller role. Building on this, I examined the relationship between leverage and employment by plotting detrended employment against detrended real household debt. The results suggest that leverage tends to follow employment fluctuations. Europe exhibits substantially greater institutional heterogeneity than the United States, encompassing variation in macroeconomic regimes, regulatory frameworks, industry exposure, cultural traditions, fiscal constraints, and labor market institutions. This dimension of cross-sectional diversity enables the identification of mechanisms in scenarios not observed in US aggregate datasets, including prolonged recessions with overextended government balance sheets. By exploiting this variety, the analysis yields external validity beyond the US's single country panels.

In this paper, I will start with a simple overlapping generations model (OLG) to show how changes in youth unemployment affect loan demand. The OLG model builds on the endowment framework of (Eggertsson, Mehrotra, and Robbins 2019), but now incorporates youth unemployment. The OLG has 4 types of agents: the employed young, the unemployed young, the employed middle-aged, and the retired old-aged. This will be an endowment economy in which middle-aged incomes will exceed those of the employed young and the retired old, providing a motive for the young to borrow from the middle-aged and the middle-aged to save for retirement, better matching borrowing and lending dynamics. The unemployed young will have less income than the employed young and will not be allowed to borrow. For the employed youth, the borrowing constraint is an external solution, meaning young agents maximize utility by borrowing up to their borrowing limit. A comparative static analysis will be conducted by exogenously increasing youth unemployment. This shows a decrease in loan demand, highlighting the significance of the borrowing constraint on the unemployed.

A HA life cycle framework with age-dependent unemployment distributions is developed to better understand borrowing and saving dynamics by capturing real constraints. When studying the comparative statics of a negative productivity shock, the firm reduces effective labor by lowering the number of households it employs. This leads to a higher unemployment rate across all age groups, especially among younger households. Similar to the OLG, this lowers loan demand but also increases precautionary savings, as working-age households with roughly zero assets begin to save more to self-insure against a higher likelihood of unemployment. This model can be studied as an impulse response of a productivity shock with different levels of persistence. This model is then compared with the same life cycle model that operates on the intensive margin and with a Huggett model. The extensive margin model shows a sharper drop in interest rates and a deleveraging in debt that becomes more apparent the more persistent the shock is. By grounding a model that uses the extensive margin in realistic demographics and unemployment realism, it captures dynamics not seen in other models.

The theoretical implications are empirically validated using Vector Autoregression (VAR) and Local Projection (LP) by examining the impulse responses of weekly hours worked, the unemployment rate, and real household debt. Both methods confirm that weekly hours worked do not change significantly, but unemployment increases, and real household debt deleverages. From there, the countries are split based on high vs. low levels of variation in weekly hours worked, unemployment, real household debt, and the household debt-to-GDP ratio. The most impactful splits were high vs low real debt volatility and high vs low average household debt, because the split with a more persistent labor productivity shock experienced larger and more prolonged increases in unemployment and a larger decrease in real household debt. This aligned with the persistence of the developed HA model.

The findings suggest that traditional monetary policy tools may be less effective at resolving prolonged household deleveraging episodes when labor productivity shocks push younger cohorts out of work. This is because employment based borrowing constraints drive persistent debt drawdowns. For policymakers and modelers alike, incorporating extensive margin dynamics into structural frameworks with agent life cycles is essential to better anticipate credit crunches and design interventions.

2 Literature Review

Macroeconomic modeling has undergone significant development since the introduction of the Real Business Cycle (RBC) framework by (Kydland and Prescott 1982) used a setting of fully flexible prices and rational expectations to attribute economic fluctuations primarily to technology shocks. However, this framework offered a limited account of labor market dynamics, particularly by the representative agent's hours worked as the variation in total hours worked (intensive margin). (Heckman 1984) pointed out that the RBC doesn't take into account the fact that a majority of business cycle fluctuations go through the extensive margin.

(Hansen 1985) extended the RBC model by incorporating the extensive margin of labor by modeling employment fluctuations as changes in the number of individuals working rather than just hours per worker. The indivisible labor approach aligned more closely with the observed volatility in employment data. New Keynesian (NK) models gained prominence by introducing nominal rigidities and emphasizing the role of monetary policy. (Rotemberg and Woodford 1997) introduced price stickiness into a DSGE structure, while (Clarida, Gali, and Gertler 1999) advanced this framework to analyze interest rate rules and inflation dynamics with the three equations of NK models. Building on this foundation, (Galí, Smets, and Wouters 2012) incorporated the extensive labor margin into New Keynesian models, improving their ability to capture the cyclical behavior of unemployment. These advances marked an important evolution in macroeconomic theory, integrating richer labor market structures and frictions into models used for policy analysis.

The extensive margin is studied not only in macroeconomics, but also in microeconomics. The microeconomic literature tends to focus on the extensive margin of labor supply. While this paper's model doesn't have shifts to labor supply, the literature consistently supports the conclusion that the extensive margin, changes in employment status, is the primary channel through which labor market adjustments occur. The extensive margin of labor supply is the decision of whether to participate in the labor force and has been a central focus in microeconomic research since the foundational work of (Heckman and MaCurdy 1980). Their analysis emphasized the importance of distinguishing between participation (extensive margin) and hours worked (intensive margin) in a heterogeneous agent life cycle model of women. This contribution laid the groundwork for subsequent empirical studies, such as (Meyer 2002), who showed how policy tools like the Earned Income Tax Credit significantly affect participation decisions among low-income individuals. More recently, (Attanasio, Levell, Low, and Sánchez-Marcos 2015) and (Chetty, Guren, Manoli, and Weber 2013) have estimated labor supply elasticities at the extensive margin, finding that the extensive margin plays a dominant role in explaining overall labor supply responsiveness, particularly for women and lower-income groups. These papers provide microfoundations for the role of the extensive margin in labor market dynamics.

The foundation of heterogeneous agent (HA) macroeconomics was laid by a series of seminal works addressing individual income risk, borrowing constraints, and precautionary saving. (Bewley 1986) developed the theoretical class of models in which agents face uninsurable idiosyncratic income shocks and engage in self-insurance through precautionary saving of government bonds, which formed the basis of incomplete markets in macroeconomics. (Deaton 1991) built on precautionary saving behavior in such an environment, showing how liquidity constraints and income uncertainty drive consumption dynamics. (Imrohoroğlu 1989) quantified the welfare losses due to incomplete markets, highlighting the inefficiencies arising from a lack of full risk-sharing. (Huggett 1993) allowed the borrowing constraint to be negative. This choice meant that agents could self-insure by saving and by borrowing against future income. (Aiyagari 1994) extended these insights to a general equilibrium setting with capital, demonstrating how precautionary savings affect aggregate outcomes and wealth inequality. (Krusell and Smith 1998) pushed the literature forward by solving a model with idiosyncratic and aggregate risk, showing that a small number of state variables can approximate the distributional dynamics effectively. More recently, the HA literature has moved towards continuous time models and away from the discrete time models that were previously used. (Kaplan and Violante 2014) introduced a two-asset HA model to explain high marginal propensities to consume among ”wealthy hand-to-mouth” households. (Kaplan and Violante 2018) review how incorporating microeconomic heterogeneity, such as differences in income and wealth, into macroeconomic models significantly improves our understanding of how households react to aggregate shocks. (Auclert 2019) provided a unifying framework to decompose the transmission of monetary policy into several redistribution channels in a one-asset HANK model, emphasizing how heterogeneity amplifies aggregate demand effects. Together, these papers represent the evolution of HA macro from theoretical foundations to empirically rich models of business cycles and policy transmission.

Heterogeneous agent models have emphasized the role of the extensive margin and how it amplifies macroeconomic fluctuations, particularly through the precautionary savings motive. (Krueger, Mitman, and Perri 2016) incorporate an endogenous extensive margin that changes labor transition dynamics based on the state of the economy into a quantitative model of the U.S. labor market, showing that unemployment risk and borrowing constraints substantially affect consumption and savings decisions. (Rendahl 2016) examines fiscal policy during periods of deep unemployment, demonstrating that government spending is particularly effective when households face elevated job uncertainty. (Challe, Matheron, Ragot, and Rubio-Ramirez 2017) build a HANK model where employment status is endogenous, and aggregate demand is driven by fluctuations in precautionary savings: when uncertainty rises, agents increase precautionary savings, which lowers real interest rates, deepening recessions. (McKay and Reis 2016) highlight the stabilizing role of automatic fiscal stabilizers, like unemployment insurance, within a HANK framework where labor market frictions and the extensive margin have a critical role in dampening shocks. (Ravn and Sterk 2017) further develop this narrative by showing how job uncertainty can endogenously generate demand-driven recessions. Collectively, these papers show that the interplay between the extensive margin unemployment and precautionary motives is central to understanding macroeconomic dynamics in heterogeneous agent models.

Life cycle dynamics serve as a crucial foundation for realistic macroeconomic analysis because they capture how households' consumption and saving behavior systematically evolve with age, driven by heterogeneous precautionary motives across life cycle stages. The work (Gourinchas and Parker 2002) aligns conventional macroeconomic models with observed microdata by establishing how a life cycle model comprehensively explains early-career liquid asset management and later-life accumulation contingent upon expected income pathways. (Eggertsson, Mehrotra, and Robbins 2019) used overlapping generations to study secular stagnation and identify effective policy response. They used demographic changes such as aging populations to capture the "secular stagnation" phenomenon more realistically than borrowers and savers models. (Kaplan, Mitman, and Violante 2020) establish that shifts in future housing demand beliefs are the principal determinant of house price appreciation, more so than credit market conditions. A life cycle model was used so to match wealth profiles and consumption saving behaviors that change significantly over a households lifespan. Life cycle modeling provides essential realism to macroeconomic frameworks by capturing how household consumption and saving behaviors evolve across different life stages.

While existing models have provided valuable insights, they often miss key features that limit their ability to match empirical labor market dynamics and household finance data. Many are formulated in discrete time, which introduces rigidity and fails to capture the nuanced timing of employment transitions as established in (Shimer 2012). Furthermore, these models frequently impose restrictive borrowing assumptions, often denying non-zero asset positions, which prevents them from realistically capturing asset accumulation, asset distribution, and consumption smoothing behavior. Moreover, many lack a fully specified life cycle framework with empirically grounded borrowing constraints, limiting their ability to reflect the economic decisions individuals make at different stages of life and based on their past employment paths. These omissions constrain the models' ability to match observed labor dynamics and household leverage patterns. This paper addresses these gaps by introducing a heterogeneous agent life cycle model with borrowing and employment dynamics that more accurately reflect real-world constraints and behaviors. These advancements contribute to the literature.

3 Beyond Aggregate Data: Exploring labor market variation and debt

3.1 Overview

The Purpose of this section is primarily to motivate the extensive margin and its correlation with household debt. It is well known that the US primarily experiences labor fluctuation over the extensive margin. The purpose of this section is to first confirm that this trend continues in Europe before examining the relationship between real household debt and employment. Looking at Europe is beneficial for a few things. First, by looking at 24 European countries, there are more observations because there is high-quality quarterly data on the country level. There is limited data on quarterly household debt per state, as well as state labor markets. This means the US could only be looked at in aggregate and lacks granularity. Second, Europe is used because it provides greater cross-sectional variation in labor market risk than is available in aggregate U.S.

data. During the period of study, these European countries had different average unemployment levels, employment fluctuations, labor market flexibility, social insurance systems, government expenditures, and so on. This creates a richer panel than using only aggregate U.S. data.

3.2 Data

To address the extensive margin, I used Eurostat employment and weekly hours worked data. This included Austria, Belgium, Bulgaria, Croatia, Denmark, Finland, France, Germany, Greece, Hungary, Ireland, Italy, Luxembourg, Netherlands, Norway, Poland, Portugal, Romania, Slovenia, Spain, Sweden, and Switzerland. It is important to note that while there was data on Cyprus, Estonia, Iceland, Latvia, Lithuania, and Malta, they weren't included because of either small size or closeness to Russia during the period, and the possible effects of the Russian and Ukrainian conflict. The UK was also not in the data, as Eurostat data didn't have weekly hours worked. To look at the correlation of employment, the extensive margin, and real household debt, we use BIS data for the household debt and HIPC price data from Eurostat. This removes Belgium, Croatia, Romania, and Slovenia as they were not part of the BIS household debt data. This dataset is well-suited to the question as it provides consistent cross-country data.

The data will first be used to decompose the variance in total hours worked into the percent deviation caused by detrended employment (extensive margin), detrended weekly hours worked (intensive margin), and the covariance between the two. This will highlight the significance of the extensive margin. The data will also be used to examine the correlation between detrended employment and detrended real household debt. The decision to focus on Eurostat/European data is deliberate. There is high-quality data for a variety of countries that have faced different shocks. This provides greater granularity into trends across different countries and can be aggregated into the overall data. Later on in this paper, when we look at VAR and LPs we will be able to split the counties into groups based on weekly hour worked variance, unemployment rate variance, real household debt variance, and household debt to GDP. This setup allows for a clearer identification of how employment outcomes relate to household borrowing behavior through different IRF splits.

3.3 Empirical Methodology

The variance decomposition of total hours worked follows the methodology outlined in (Hansen 1985). To extract cyclical components, all series are detrended using the Hodrick-Prescott (HP) filter with a smoothing parameter of 1600, which is appropriate for quarterly data. The first two and last two years of data for each country was thrown out due to the poor fit of tails when using HP filters. The detrended value is then calculated as the ratio of the detrended (cyclical) component to the trend, plus one, a formulation that facilitates interpretation in log-linear terms. The variances and covariances are computed using the natural log of the detrended value of each term. It is important to note that covariances can be negative if the underlying variables are negatively correlated. The variance decomposition following equation (1). T stands for total hours worked in an economy over the past year. E stands for employment. H stands for weekly hours worked. The decomposition is performed in an equal-weighted manner across all data points to construct aggregate measures. Additionally, a within-country variance decomposition is conducted for each country in the sample.

$$\mathrm{var}(\ln T) = \mathrm{var}(\ln E) + \mathrm{var}(\ln H) + 2\,\mathrm{cov}(\ln E, \ln H) \tag{1}$$

To examine the relationship between labor market dynamics and household leverage, detrended employment and detrended real household debt were analyzed. It is important to note that real household debt was detrended the same way employment was. A one quarter lag in detrended employment was selected to maximize the empirical fit, reflecting the realistic delay with which households respond to labor market conditions by adjusting their borrowing behavior, both in leveraging during periods of high employment and deleveraging during periods of lower employment. The correlation is then conducted in an equal-weighted fashion across all available quarterly-country observations to construct aggregate patterns. Within-country heterogeneity was assessed to capture the heterogeneity in the employment and household debt correlation within each country.

3.4 Falsification Tests

The aggregate variance decomposition was computed in an equal-weighted manner across all data points to ensure that no single country disproportionately influences the aggregate results. This provides a more balanced view of overall labor market dynamics within Europe, as different countries dealt with different idiosyncratic shocks. This approach allows for the construction of aggregate measures that reflect the central tendencies of the countries as a whole, rather than being skewed by larger economies. Additionally, within-country decompositions were done to show important heterogeneity across nations, enabling the identification of distinct labor market patterns that may be masked in the aggregate. Together, these two levels of analysis reinforce the robustness of the findings by showing that the significance of the extensive margin is not driven by outliers and holds consistently across multiple contexts. To strengthen my analysis of the relationship between employment and real household debt, I first examine detrended employment and detrended real household debt at the country level. It indicates a correlation, and the choice to use lagged employment deviation. Specifically, a quarter lag aligns empirically and intuitively with the time it takes for households to adjust their borrowing behavior in response to labor market changes.

Once the country-level patterns are established, an aggregate correlation is constructed by pooling all data points with equal weighing, ensuring no single economy skews the results. This dual approach, both within the country and in the aggregate, strengthens the robustness of the findings by capturing both aggregate and within-country dynamics.

3.5 Empirical Results

In Figure (35), the aggregate variance decomposition of total hours worked shows that fluctuations in the employment rate account for over 83.2 % of the variation across countries. The second-largest contributor is the variance in weekly hours, at 18.2 %. The covariance term is negative and near zero, indicating a near-zero negative correlation between employment and weekly hours worked. This analysis underscores the pivotal role of the extensive margin, driven by fluctuations in employment, in explaining variations in total hours worked across European countries.

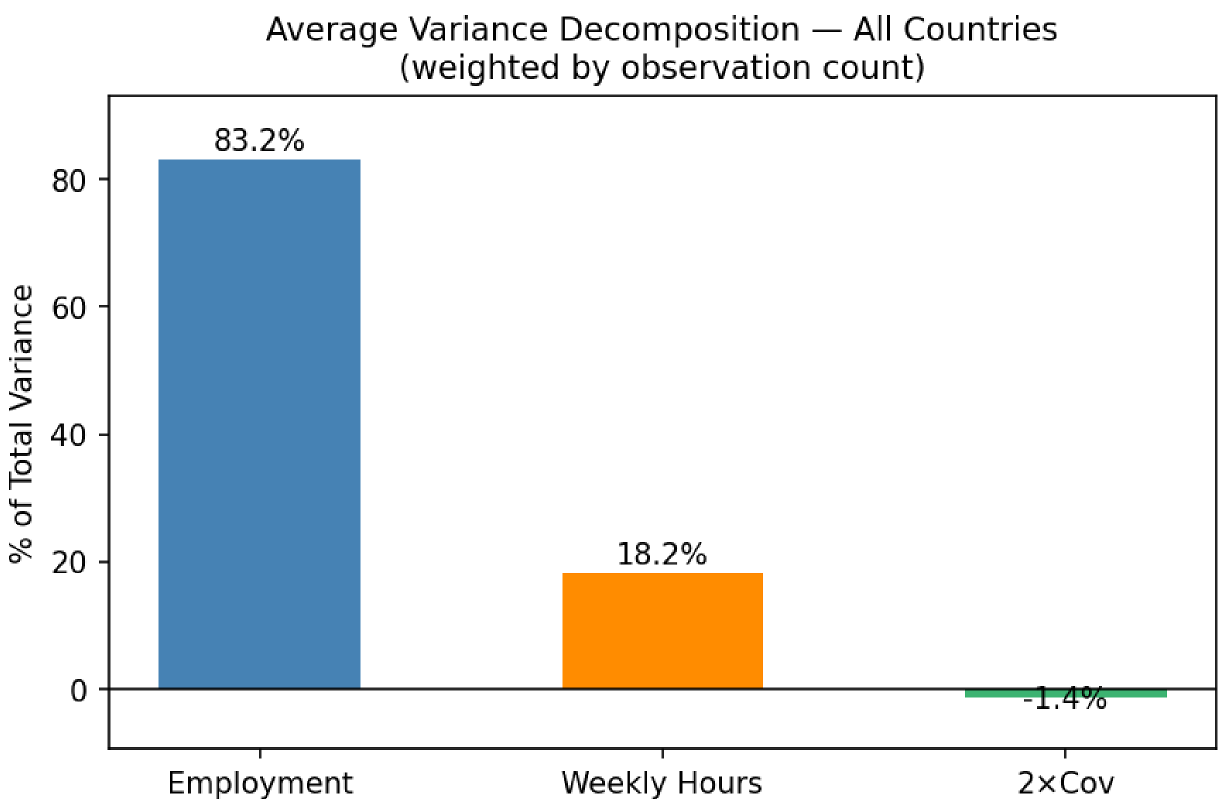

Figure 1: Aggregate Variance Decomposition

Figure (36) breaks down the variation per country. There are a few things that stand out. First, some countries: Austria, Denmark, Czechia, France, the Netherlands, Portugal, Romania, Slovenia, and Switzerland show more than 100 % share in total variance. This means that the total variance caused by employment and weekly hours worked is greater than the total variance, and the covariance term is negative because employment and weekly hours worked are negatively correlated. The green bar going back to 100 % shows a negative contribution of the covariance term. Second, Austria is the only country that has a higher percentage of total variance being caused by weekly hours worked. This is due to Austria's relatively stable employment rate during the Great Recession, which only rose modestly from 4.4% in 2008 to 5.7% in 2009, and remained below 5.7% until 2014. This reflected a milder recessionary impact on the employment rate.

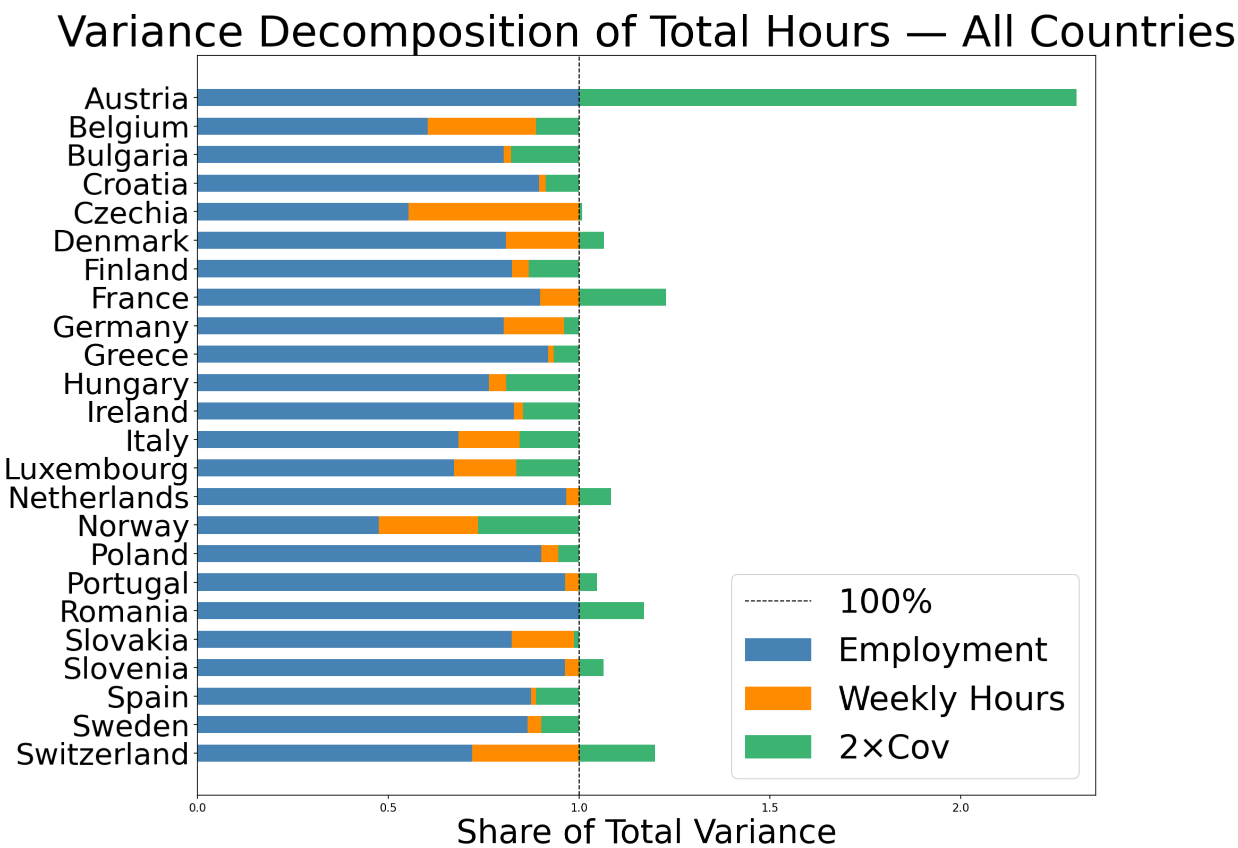

Figure 2: Within Country Variance Decomposition

Figure (3) illustrates the relationship between detrended real household debt and detrended employment. The two series generally display strong positive comovement, suggesting a meaningful correlation between labor market conditions and household borrowing behavior. The correlation is the strongest when employment is lagged one quarter/period to account for the time needed to increase and decrease leverage. Household debt dynamics are closely tied to labor market developments, which supports the use of the employment rate (extensive margin) as a key contributor to an economy's credit dynamics.

Figure 3: Detrended Household Debt and Employment

Figure (4) presents the aggregate relationship between detrended real household debt and the detrended employment rate lagged one quarter. The figure shows a strong positive correlation, indicating that periods of higher employment are associated with higher household leverage, while periods of lower employment are associated with lower household leverage. This aggregate relationship provides further evidence that household debt dynamics are systematically linked to the extensive margin of labor fluctuations.

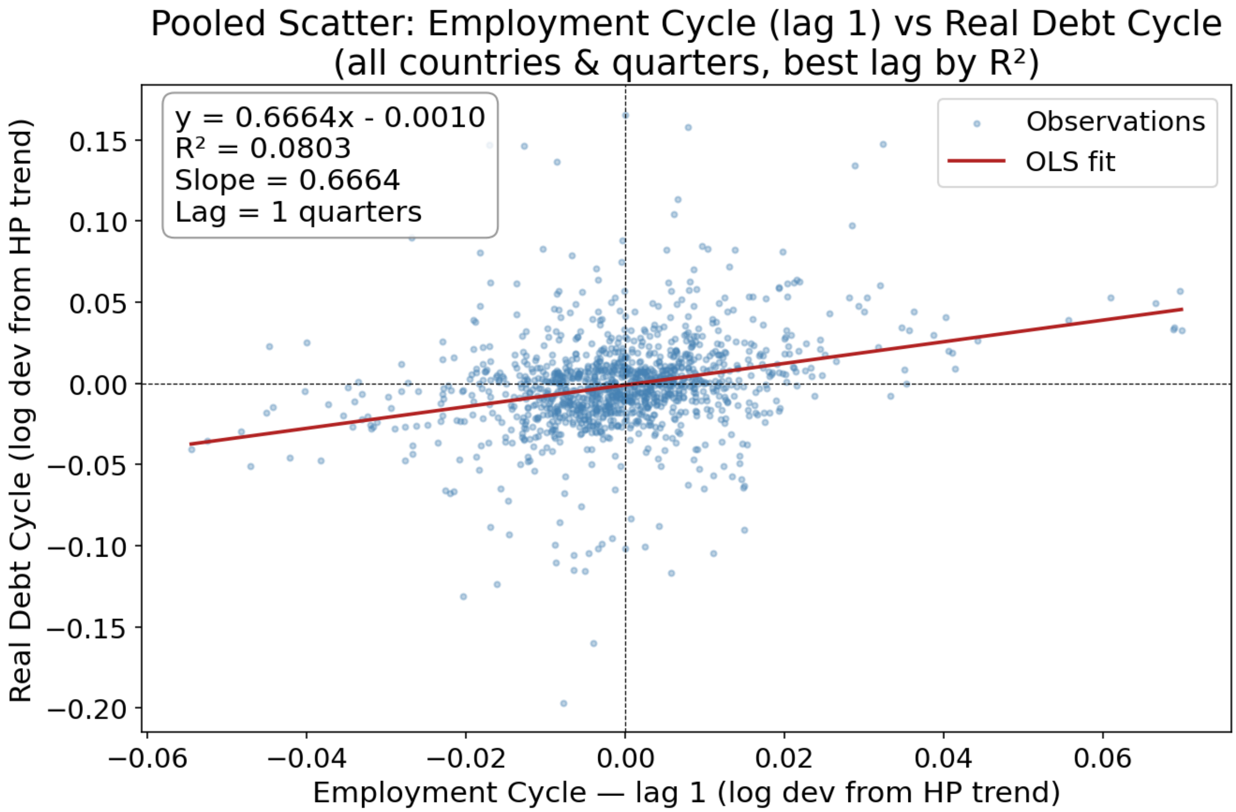

Figure 4: Aggregate Household Debt and Employment

A within-country analysis was conducted, Figure (5), to verify whether the positive relationship between lagged detrended employment and real household debt holds across individual countries. Almost all countries exhibit a positive slope, indicating a positive correlation between the two variables. Notably, Germany, Hungary, Luxembourg, and the Netherlands have a negative slope or inverse relationship, with Luxembourg and the Netherlands having an almost zero slope. This means that 16 of the 20 countries have a positive correlation. It is important to note that Czechia has a very strong relationship that could be skewed by outlier points with lower cyclical real household debt. Overall, looking at the country-level confirms that there is a positive correlation between employment and household debt.

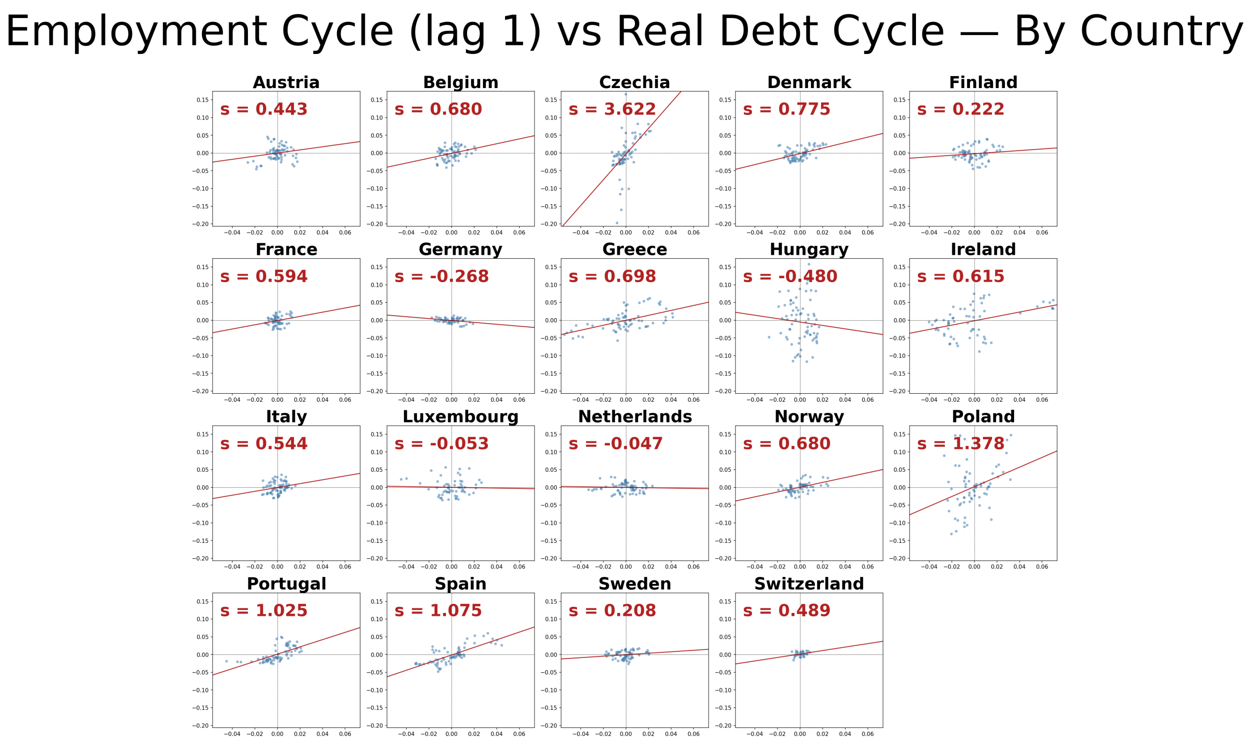

Figure 5: Within Country Household Debt and Employment

In conclusion, the empirical analysis highlights the existence and importance of the extensive margin in explaining labor market dynamics. As shown in Figure (35), the variance decomposition of total hours worked indicates that variation in employment is the dominant contributor to the total hours worked variance, surpassing both weekly hours worked and the covariance term. This finding is further supported by the within-country decompositions Figure (36), which reinforces the prominence of the extensive margin despite some country-specific deviations. Moreover, the relationship between real household debt and employment, as illustrated in Figure (3) and Figure (4), and Figure (5) highlights the significance that employment fluctuations, the extensive margin, has over household leverage. A positive correlation is observed at both the aggregate and country levels. This suggests that improvements in employment conditions are associated with increased household borrowing, whereas downturns coincide with deleveraging. These results collectively support the hypothesis that labor market variation through the extensive margin plays a key role in shaping both labor and loan dynamics.

4 Isolating the Demand Channel: A simple overlapping generations model

4.1 Overview

The purpose of the simple overlapping generation (OLG) model is to demonstrate how changes in youth unemployment, the extensive margin of youth labor, affect loan demand. Loan supply will be analyzed more thoroughly in a heterogeneous agent (HA) framework. The OLG model is used to isolate the demand channel clearly and concisely. It builds on the endowment-based OLG framework in (Eggertsson, Mehrotra, and Robbins 2019) and extends it by incorporating youth unemployment. Households live for three periods: the young (20–39), the middle-aged (40–59), and the old (60–79). There are four agent types: employed young, unemployed young, employed middle-aged, and retired old. The model does not include physical capital. The model instead focuses on intergenerational borrowing, where young agents borrow from middle-aged savers. Young households are born during the current period and consume from their low current income and borrowing, while middle-aged households use their higher earnings to repay their debt accumulated from the previous period and save for retirement. Retired households consume out of their now lower current income and accumulated savings. Borrowing is constrained by an exogenous loan-to-income limit. This borrowing constraint is assumed to be a binding external solution, so it is utility maximizing for young agents to borrow up to their limit. This feature simplifies the math and reflects realistic credit behavior. Critically, unemployed youth receive lower income than employed youth, which tightens their borrowing capacity. As youth unemployment rises, the share of young borrowers with low income increases, reducing aggregate loan demand. This mechanism highlights how labor market conditions among the young translate into macroeconomic credit fluctuations.

4.2 Setup

The household utility function eq.(2) of the young depends on the current consumption when the household is young C_t^y, middle-aged consumption C_{t+1}^m in one period from now, and old age consumption C_{t+2}^o in two periods from now. The future utility is discounted by a rate of time preference β in one time period and by β^2 in two time periods.

$$U_t^y = ln(C_t^y) + \beta ln(C_{t+1}^m) + \beta^2 ln(C_{t+2}^o) \tag{2}$$

The debt constraint, eq.(3), is given as a fraction of youth income. It is important to note that the debt limit will form an external solution as the young households will not be able to borrow their utility-maximizing borrowing quantity. Later, we will change the youth's income based on being employed or unemployed. This will tighten the debt limit D_t for the unemployed youth. This will lower loan demand and decrease the total quantity of loans. It is important to note that borrowings B_t are less than the debt limit because your debt limit includes the interest on borrowings. The amount borrowed is consumed, not the debt limit. This allows young households to consume more when the real interest rate is lower, as that lets them borrow more with the same debt constraint.

$$D_t = B_t(1+r_t) = \gamma Y_t^y \tag{3}$$

The young budget constraint eq.(4), middle budget constraint eq.(5),and the old budget constraint eq.(6) govern the consumption that the young, middle, and old households. The young households consume all of their income and borrowings. Youth borrowings can be expressed as a function of their income and the real interest rates. The middle-aged households consume out of their current income less the their debt limit from last period when they were young, plus borrowings they take out in the current period. They will have negative borrowings (savings) in the current period, so saving in the current period reduce consumption. Old-age households will consume out of old-age income less middle-aged borrowing (plus middle-aged savings).

$$C_t^y = Y_t^y + B_t^y = Y_t^y\left(1+\frac{\gamma}{1+r_t}\right) \tag{4}$$

$$C_{t+1}^m = Y_{t+1}^m - (1+r_t)B_t^y + B_{t+1}^m \tag{5}$$

$$C_{t+2}^o = Y_{t+2}^o - (1+r_{t+1})B_{t+1}^m \tag{6}$$

The Euler equation, eq.(7), is derived from eq.(2).

$$\frac{1}{C_t^m} = \beta E_t(1+r_t)\frac{1}{C_{t+1}^o} \tag{7}$$

By combining the middle-aged budget constraint eq.(5) and the old-age budget constraint eq.(6), and applying the Euler equation eq.(7), we derive the consumption function for middle-aged agents, shown in eq.(8).

$$C_{t+1}^m = \left(\frac{1}{1+\beta}\right)\left(Y_{t+1}^m + \frac{Y_{t+2}^o}{1+r_{t+1}} - (1+r_t)B_t^y\right) \tag{8}$$

4.3 Adding Youth Unemployment

To add richness to the model, the young are divided into two employment states: employed and unemployed. This split captures the extensive margin of employment, with ω_t percent of the young employed and 1-ω_t unemployed. This differs from the empirical section: instead of focusing on the number of employed individuals, we focus on the unemployment rate because it indicates economic distress and scales across economies. It is assumed that middle-aged households are fully employed. This assumption reflects the observation that youth unemployment tends to be more volatile than overall unemployment.

Young individuals disproportionately bear the brunt of labor market fluctuations for two main reasons: they are more susceptible to layoffs because of their shorter tenure, and they are more affected by hiring slowdowns because they typically enter the labor market from unemployment. (Shimer 2012) highlighted that the primary driver of rising unemployment is reduced hiring rather than increased layoffs. For simplicity, the model assumes that all of the middle-aged are employed. The focus is placed on the extensive margin of youth unemployment, as the young are the primary drivers of loan demand both empirically and in OLG literature. The unemployed youth earn 1/3 of the employed youth's income. The old-age receive social security income. Old age income lowers loan supply and increases real interest rates, as this reduces the percent of income that needs to be smoothed for consumption next period when the agent is old-aged. This setup demonstrates that fluctuations in the real interest rate are primarily driven by changes in youth unemployment, consistent with the empirical evidence.

The exogenous budget constraint eq.(9) is now endogenously determined by the level of youth unemployment. This impacts loan demand because the unemployed cannot borrow as much: their debt constraint, D^u, is zero and significantly lower than the employed's, D^e. The extensive margin of young unemployment now directly lowers loan demand due to the lower debt limit. The equation for loan demand is now the fraction of employed youth times the employed debt constraint plus the fraction of unemployed youth times the unemployed debt constraint divided by the gross real interest rate. This equation shows that increasing the unemployment rate reduces loan demand.

$$L^d = \frac{\omega_t D^e + (1-\omega_t)D^u}{1+r_t} \tag{9}$$

Loan supply, on the other hand, is an interior solution determined by the Euler equation (eq.).(7) and the middle-aged budget constraint eq.(5) . Loan supply is positively affected by middle-aged income and real interest rates. Loan supply is negatively affected by income in old age and by the debt the middle-aged accumulated in the last period when they were young. This implies that households that experienced unemployment when they were young accumulated less debt and smoothed their higher wealth into old age consumption. Since income during the middle-aged period is exogenously fixed and independent of earlier employment history, these households adjust their intertemporal consumption accordingly. However, this behavior is not the model's central focus; it is primarily designed to illustrate loan demand dynamics on the extensive margin of youth employment.

$$L^s = \left(\frac{\beta}{1+\beta}\right)(Y_t^m - \omega_{t-1}D_{t-1}^e - (1-\omega_{t-1})D_{t-1}^u) - \left(\frac{1}{1+\beta}\right)\frac{Y_{t+1}^o}{1+r_t} \tag{10}$$

4.4 Results

Loan demand and loan supply are shown in Figure (6). The youth unemployment in the steady state is assumed to be 5%. The youth unemployment rate is assumed to be the same in the current and past periods. The figure shows that loan demand is much more inelastic than labor supply. When the youth unemployment rate increases, the red line shifts to the magenta line, and the loan quantity decreases more than the real interest rate. Conversely, when the middle age group was unemployed in their youth, there was a small increase in the loans supplied as we move from the blue line to the light blue line. This means that the impact of youth unemployment on loan demand has a greater influence on loan quantity and interest rates than changes in loan supply. Adjustments over the extensive margin, youth unemployment, play a central role in shaping aggregate loan demand, as fluctuations in youth employment status directly impact borrowing behavior, making it a critical channel through which labor market dynamics influence loan quantity and real interest rates.

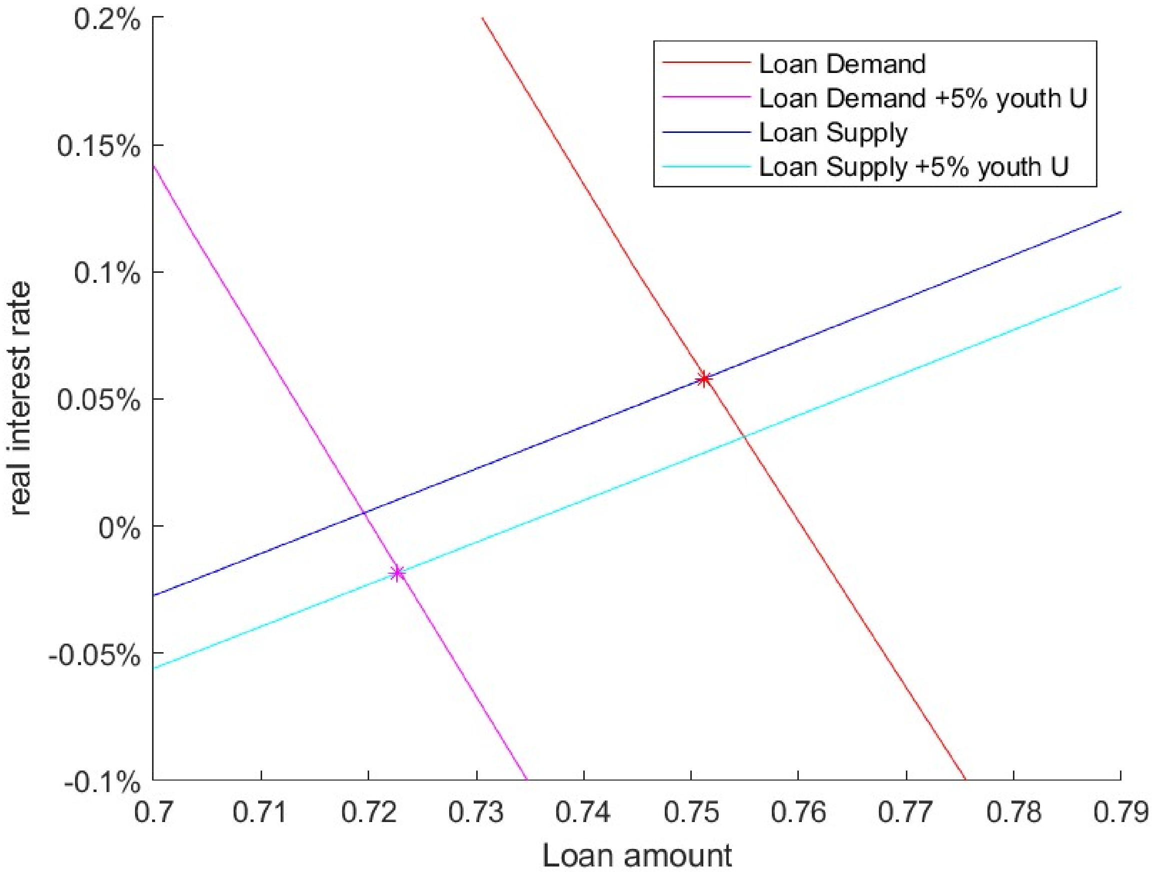

Figure 6: Loan Demand and Loan Supply

Equating loan demand eq.(9) with loan supply eq.(10) yields equilibrium real interest rate eq.(11). The real interest rate increases with higher youth employment in both the current and the previous period, as observed in Figure (7). Youth unemployment in the current period lowers the real interest rate by tightening borrowing constraints, as a greater share of credit-constrained unemployed youth reduces aggregate loan demand, exerting downward pressure on equilibrium rates. This relationship arises because unemployed youth borrow less, leading to lower indebtedness when they enter middle age. Consequently, middle-aged households that experienced youth unemployment retain a larger share of their income after debt repayment. Given their desire to smooth consumption over time, as captured by the Euler equation, these households increase their loan supply in the present period. Thus, higher past youth unemployment raises current loan supply, placing downward pressure on the equilibrium interest rate.

$$1 + r_t = \frac{(1+\beta)(\omega_t D^e + (1-\omega_t)D^u) + (\frac{1}{1+\beta})Y^o_{t+1}}{\beta(Y^m_t - \omega_{t-1}D^e_{t-1} - (1-\omega_{t-1})D^u_{t-1})} \tag{11}$$

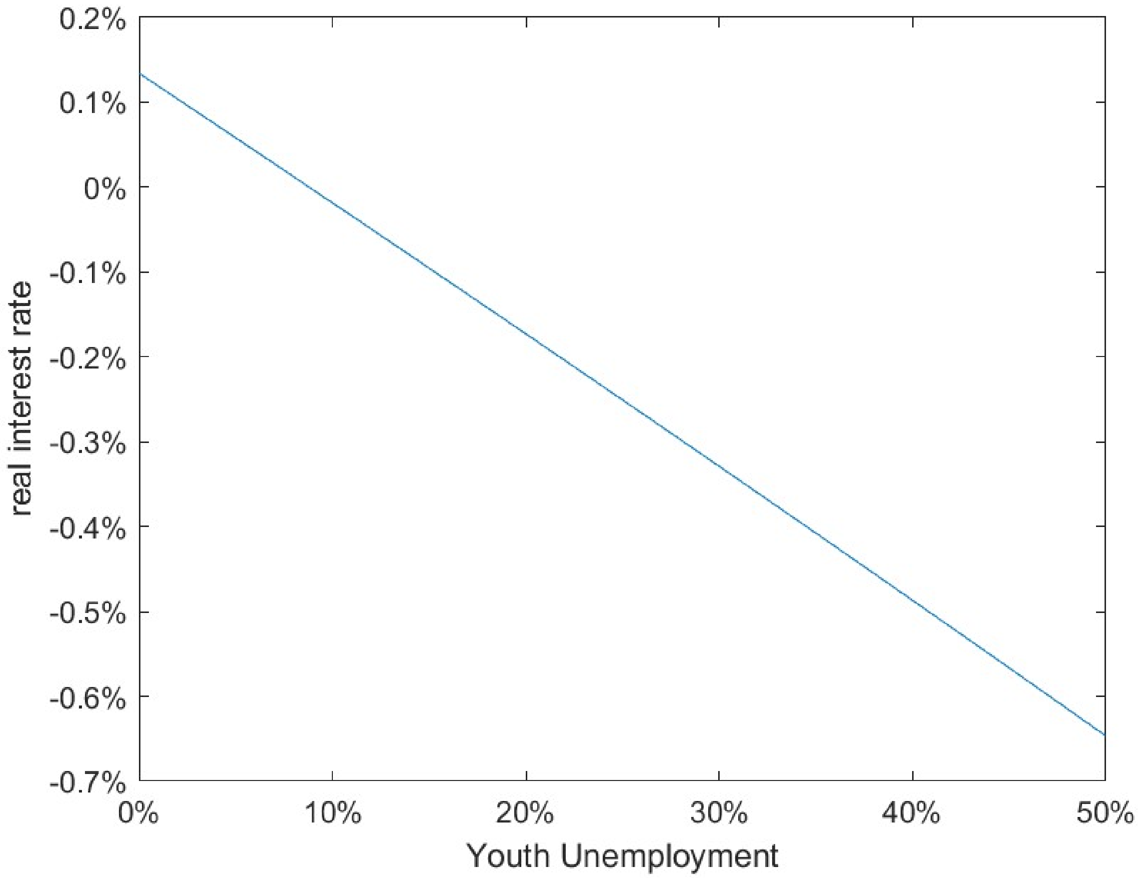

Figure 7: Real Interest Rate based on Youth Unemployment

Figure (8) illustrates the total loan quantity as a function of steady-state youth unemployment. As youth unemployment increases, the total volume of loans in the economy declines.

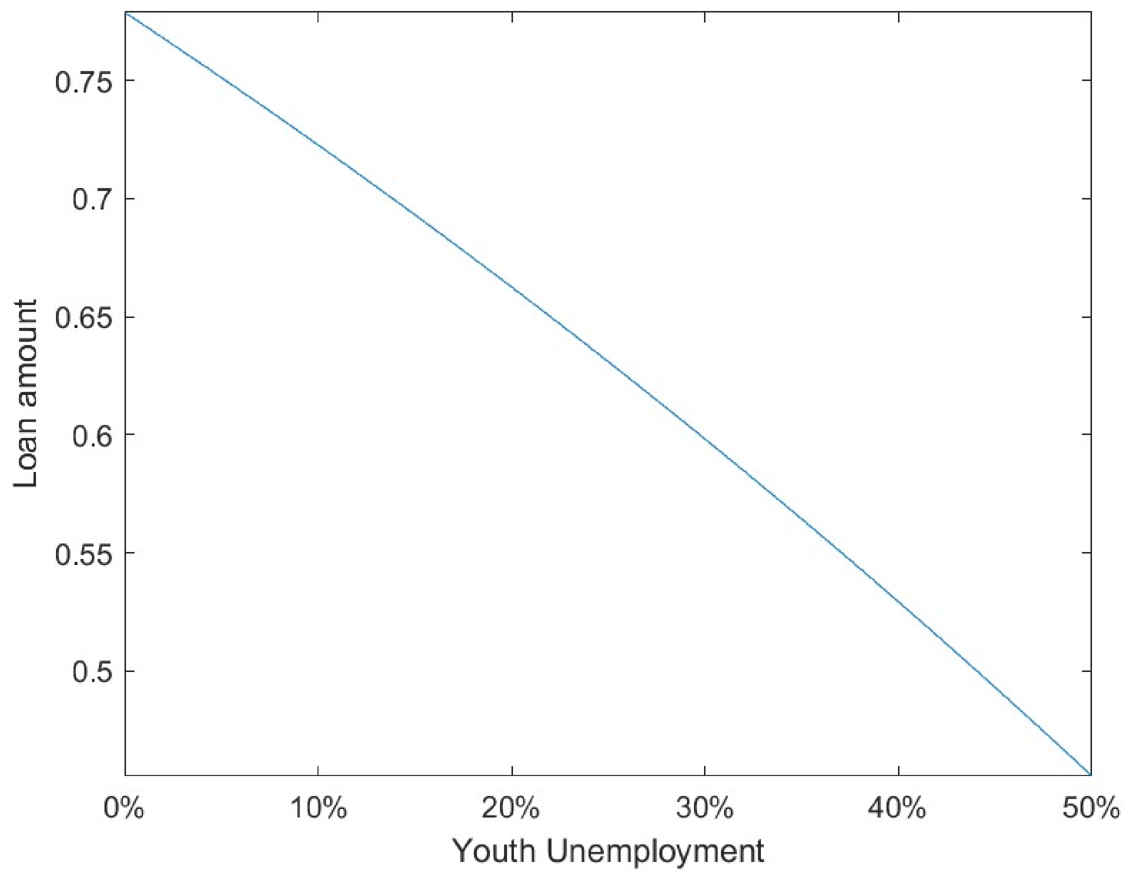

Figure 8: Loan Quantity based on Youth Unemployment

In conclusion, this paper develops an OLG model to analyze how the youth unemployment, the extensive margin, influences loan demand, loan supply, and the aggregate outcomes of interest rates and the quantity of debt outstanding. By introducing unemployment heterogeneity among young households, the model captures how borrowing constraints and income differences across employment states affect intertemporal consumption and savings behavior. The results demonstrate that youth unemployment significantly depresses aggregate loan demand by tightening borrowing constraints on the unemployed. Loan supply increases modestly when middle-aged households experience unemployment in the previous period, as this reduces prior borrowing. The dominant effect arises on the demand side, and loan supply is not a focus of this model; it is just a mechanical feature that needs to be explained. The real interest rate is shown to fall with higher youth employment, reflecting the dominant shift in loan demand caused by the borrowing constraint. Overall, the analysis highlights the central role that fluctuations in youth employment play in shaping credit market dynamics and underscores the importance of the extensive margin in understanding fluctuations in loan demand and their effects on loan quantity and real interest rates.

5 Microfoundations of Interagent Lending: Incorporating extensive margin risk into the HA framework

5.1 Overview

The heterogeneous agent (HA) model is employed to examine how unemployment risk over the extensive margin with incomplete markets influences agent consumption and saving behavior. Unlike the overlapping generations (OLG) model, the HA framework better replicates real-world dynamics and more accurately reflects shifts in loan supply due to the precautionary savings motive. The precautionary savings motive was not incorporated in the OLG, only the intertemporal and consumption smoothing motivations were incorporated. This analysis explores not only how changes in loan supply and demand affect the quantity of household loans and real interest rates, but also how saving behavior differs between adjustments along the intensive and extensive margins. In particular, the distinction in saving rates for agents on the same grid location (asset level, employment status, and age) is analyzed between the two adjustment mechanisms.

5.2 Setup

5.2.1 Utility Function and Budget Constraint

The model has a standard Constant Relative Risk-Aversion (CRRA) utility curve eq. (12). There is no disutility for work. This assumption will make all households under the retirement age want to work regardless of asset level, and no agents above retirement age will work. ,

$$U_t = \frac{c_t^{1-\sigma}}{1-\sigma} \tag{12}$$

Eq. (13) represents the budget constraint. wz_lh_e is wage income earned. Wage income earned is based on the wage rate w times z_l the age productivity profile times hours worked h_e. The shape of income is parabolic, peaking when agents are in their mid-50s. For the steady state example, hours worked will be zero or one. Unemployment/social security u_e is given to all the unemployed individuals, including the retired. π_l is the dividend income from firms. The profile of the dividends matches the z_l income profile. Interest payments are assumed to be untaxed in order to isolate the real rate of return. Taxing interest income would introduce a wedge between the gross and net return, affecting the equilibrium real interest rate. Since the focus of this analysis is on the real return absent tax distortions, taxation is omitted from interest income.

$$c(a,e,l) = (wz_lh_e + u_e(1-e) + \pi_l)(1-\tau) + ra - s(a,e,l) \tag{13}$$

Eq.(14) is the recursive value function known as the Hamilton-Jacobi-Bellman (HJB). The rate of time preference is 4% in line with the findings from (Gourinchas and Parker 2002). The solution to the HJB gives a value function $v(a,e,l)$, a policy function for consumption $c(a,e,l)$, and a policy function for savings $s(a,e,l)$ for each discrete agent based on asset level, employment status, and age. Life cycle and employment flows are exogenous to the value function, and the 3 policy functions are endogenous.

$$\rho v(a,e,l) = \max_{c} \; u(c) + v_a'(a)\dot{a} + \sum_{e' \neq e} \lambda_{ee'}(v(a,e',l) - v(a,e,l)) + \sum_{l' \neq l} \lambda_{ll'}(v(a,e,l') - v(a,e,l)) \tag{14}$$

The first-order condition of the HJB, eq.(15), states that the marginal utility of consumption will equal the partial derivative of the value function with respect to assets. The FOC is necessary to convert the HJB into a viscosity solution to allow numerical methods to be performed.

$$u^{'}(c) = v_a^{'}(a,e,l) \tag{15}$$

To solve the function computationally, the inverse of the marginal utility function of the derivative of the asset grid $u^{'-1}(v_a^{'}(a,e,l))$ is substituted into where consumption c was. This yields the viscosity solution eq.(16). This eliminates consumption from the equation and allows the equation to be solved recursively based on grid location and exogenous transition dynamics.

$$\begin{aligned} \rho v(a,e,l) = \max_{c} \; u(u^{'-1}(v_a^{'}(a,e,l))) + \\ v_a^{'}(a)((wz_l h_e + u_e(1-e) + \pi_l)\tau_1 + ra - u^{'-1}(v_a^{'}(a,e,l))) + \\ \sum_{e' \neq e} \lambda_{ee'}(v(a,e',l) - v(a,e,l)) + \sum_{l' \neq l} \lambda_{ll'}(v(a,e,l') - v(a,e,l)) \end{aligned} \tag{16}$$

5.2.2 Asset, Employment, and Life cycle Dynamics

Income profiles are calibrated using age-earnings data from the 2013 Bureau of Labor Statistics (BLS) report. The data is fitted using ordinary least squares (OLS), with age and age squared as covariates. All employed individuals within the same age group are homogeneous, implying identical income dependent on age. While incorporating income heterogeneity within age groups would better match real-world dynamics, it would also introduce significant computational complexity at a stage where lending and borrowing dynamics are already captured by life cycle income dynamics. This specification results in a parabolic income profile over the life cycle, with earnings peaking at 50. The estimated income curve is then scaled such that aggregate productivity equals one at the natural rate of employment of the working-aged individuals. The scaling accounts for full retirement at age 65, with all individuals 65 or older being retired. By matching the shape of the life cycle income profile, the HA model better matches life cycle dynamics compared to the OLG framework. The model retains the standard borrowing motive for the young and a retirement savings motive for the middle-aged. Consistent with (Guerrieri and Lorenzoni 2017), unemployment benefits and Social Security income are both modeled as 40 percent of average labor income. Dividend income is modeled as age-dependent and follows the same profile as labor income. It is assumed that both employed and unemployed individuals receive the same dividend income. Dividend income is scaled differently, as it is distributed across all age groups, not just the working population.

Agents enter the workforce at age 20, and exit the model upon death at age 75. The life cycle is discretized into 11 five-year age groups: 20, 25, 30, 35, 40, 45, 50, 55, 60, 65, and 70. Each agent is assigned the lower bound of their age group; for example, an individual in the 20–24 group is treated as age 20. Because the model is in continuous time, agents age according to a Poisson process. Each quarter, an agent has a probability of $\frac{1}{20}$ of transitioning to the next age group or dying if the agent is 70. The distribution of agents across age is assumed to be uniform. For every agent that reaches age 75 and exits the model, a new agent is introduced at age 20. New entrants inherit the agent's asset holdings upon death. This includes if the exiting agent holds a negative asset position. Due to the stochastic nature of the Poisson aging process, some agents may age faster than they pay down their debt, as they may be randomly selected to advance in age without a corresponding opportunity to pay down debt. Although rare, this setup allows dying agents to hold negative assets.

The borrowing constraint is modeled using a maximum payment-to-income (PTI) ratio eq.(17). Following the methodology of (Kaplan, Mitman, and Violante 2020), the minimum loan payment is set so the loan fully amortizes by the time the agent reaches retirement age. The PTI limit is set to 28% of current income, consistent with the front-end PTI ratio documented in (Quercia, McCarthy, and Wachter 2003). The borrowing constraint is further enhanced by restricting borrowing to employed households only. While this assumption may initially appear overly restrictive, real-world evidence shows that many households have limited consumption-smoothing capabilities. According to a Forbes article (McGrath 2016), "63% of Americans Don't Have Enough Savings to Cover a $500 Emergency," indicating that a significant portion of the population lacks adequate financial buffers. This means that while individuals are in debt, a good portion do not have adequate savings in their savings and checking accounts. Furthermore, research from (Haughwout et al. 2024) highlights that approximately 18% of credit cards are utilized up to 90% of their limits. These findings suggest that individuals with negative or near-zero net assets face substantial limitations in accessing credit and therefore have limited means to smooth consumption during periods of financial stress or unemployment. These findings provide empirical support for the assumption that there is a single asset spectrum, rather than two distinct spectrums for debt and savings.

$$pmt_t = .28 y_{e,l} e_e \tag{17}$$

The minimum debt payment is calculated using a fixed annuity formula eq.(18), reflecting how most household debt is structured in the U.S. financial system. Because the model is solved in steady state, the interest rate is assumed to be fixed. As agents age and gradually repay their debt, the required minimum payment adjusts over time. While this implies the payment profile is not constant, maintaining a fixed payment schedule would introduce an additional dimension of complexity, requiring households to retain their payment schedule despite changes in debt or age. Such an extension would be problematic, particularly in continuous time, because agents who age up more quickly (due to the Poisson process) might retire while still holding significant debt. Although this is not ideal, modeling in continuous time is essential to accurately capture the employment transition dynamics.

$$pmt_t = \frac{-a * r}{1 - (1+r)^{65-age}} \tag{18}$$

An important assumption is that agents cannot default on their debt. This assumption is required because loans are unsecured by collateral and must be backed only by future income. If an agent is constrained and unable to borrow further, they are required to make the minimum debt payment, provided their consumption remains above a minimum threshold of 0.1. Otherwise, the payment is equal to income less the minimum consumption threshold. The imposed saving policy function sets both the consumption policy function and the value function. Regions close to the forced savings regime must incorporate this constraint into the value function because they could switch to unemployed or age up. This adjusts optimal consumption and saving behavior around the forced savings constraints. This forced savings rule ensures that nearly all agents avoid entering retirement with zero or negative net assets. It is important to emphasize that most households are not subject to this constraint unless they are unemployed. In general, agents voluntarily save above the minimum requirement to pay down debt and accumulate wealth for retirement.

5.2.3 Firm and Employment Dynamics

Firms pay real wages above the market-clearing level in accordance with efficiency wage theory. As discussed in (Katz 1986), firms may pay higher wages to reduce turnover, motivate greater worker effort, discourage shirking, deter unionization, and attract higher-ability employees. Unemployment is introduced into this model due to the wage markup. In this model, the wage is set so that, in equilibrium, the unemployment rate is 6%. The wage at steady state equilibrium is set to be (1 - α) or the labor share of a traditional model. The labor supply curve is set to equal the Frisch elasticity observed in the macroeconomy. The elasticity used is 2.84, based on macro estimates of aggregate hours in (Chetty, Guren, Manoli, and Weber 2011). This includes .54 for the intensive margin and 2.50 for the extensive margin. For simplicity, the aggregate is used across both margins. It is important to note that the authors conducted micro estimates and found aggregate hours fluctuation to be around 0.82, with 0.54 from the intensive margin and 0.28 from the extensive margin. The micro estimates were based on a distribution of reservation wages and quasi-experiments. It is important to note that this is purely a focus on labor supply choices, but in order to choose to work, there has to be a firm that hires you. This completely misses involuntary unemployment and effectively treats labor force participation as the sole extensive margin. In my model, the labor supply curve is the wage based on the quantity of labor used, not the other way around. This uses empirical data to create the wage line based on the quantity of labor rather than a Nash bargaining approach. Nash bargaining can get very complex in an environment with such a large distribution of agents across age, wealth, and employment status. A different way of looking at it is that the microeconomics show a .54 Frisch elasticity from the intensive margin, a .28 from the extensive margin of a larger labor force, and a 2.84 - .54 - .28 = 2.02 from employment. I want to reiterate that this follows the path of inverse Frisch elasticity, where the wage is determined by the quantity of labor used. The intertemporal Frisch elasticities are likely understated as well because firms face no labor adjustment costs. If there is one time period when workers are willing to accept lower wages, not all of them will be hired, because firms now face a penalty for hiring the marginal agent now and laying them off later when the required wage drops.

The firm's production function is modeled as a neoclassical function with diminishing returns to labor and no capital input. Costs are real wage income, denoted as total effective labor times the wage. Together, the profit function is the production function less the real labor cost as seen in eq.(20). Total effective labor is computed in eq.(19) as the aggregate sum of employed individuals, each weighted by their age-specific productivity z_l and individual hours worked $h_{e,t}$. It is important to note that in calibration, age-specific productivity is adjusted evenly across all age groups so that at the natural rate of employment, total labor used equals 1. In the steady state, hours worked are treated as binary because individuals either work and provide one full unit of labor or zero units of labor. The firm can adjust total labor input along two margins: the extensive margin, by varying the number of employed agents, and the intensive margin, by adjusting the hours worked by employed agents. This provides a flexible framework for analyzing the difference between the intensive and extensive margins when there is a productivity shock.

$$N_t = \sum_e \int_l (z_l h_{e,t}) \tag{19}$$

$$\pi_t = A_t(N_t)^{1-\alpha} - wN_t \tag{20}$$

Employment dynamics play a pivotal role in this model's framework. Rather than incorporating a full job search model, job flow parameters are calibrated based on the empirical findings of (Shimer 2012). Specifically, the model assumes that hiring rates are determined ex-ante. This excludes individuals who enter the labor force already employed, particularly those aged 20-25. This assumption reflects real-world patterns, such as college students securing employment offers before graduation and before formally entering the labor market.

To replicate observed labor market behavior, the share of young individuals who begin employed is set as a fraction of the older cohort exiting the labor force through retirement. This calibration ensures the model achieves the desired natural spread, as observed among 20-24 year olds, with a higher unemployment rate in the steady state. Unlike some models that assume uniform job separation probabilities across workers, this framework introduces age-dependent exit probabilities for two reasons. First, agents retire at age 65, which increases the exit rate for the 60–64 age group. Second, to align the age profile of unemployment rates with empirical data, the job separation rate for young workers aged 20 and 25 is increased, matching the data during recessions. This is consistent with labor market dynamics, as younger workers are more likely to be laid off due to lower tenure and job security. The higher youth job exit probability helps calibrate the model to match targeted unemployment rates across the life cycle. The overall separation rate is calibrated to ensure that the model's steady-state equilibrium matches the natural rate of unemployment.

(Shimer 2012) observed that the average monthly hire rate was 44 %. Assuming homogeneous hiring rates, this corresponds to a quarterly hire rate of 82.4 %. This means it takes the average person roughly 3.6 months to find a job once unemployed. The job separation rates are set so that the model achieves the targeted natural unemployment rate at the steady state. (Shimer 2012) finds that 77 % of fluctuations in the unemployment rate since 1948 were because of changes to the hiring rates. For simplicity, this model assumes 75 % of the fluctuations in unemployment were caused by a reduced hiring rate. To incorporate this finding into the model, for every 3% reduction in the hiring rate, the separation rate increased by 1%. While I am not looking at the hiring and separation rates in an IRF response, this is a good start to set the employment flow dynamics for different natural rates of unemployment.

5.2.4 Market Clearing Conditions

The Kolmogorov Forward Equation eq.(21) ensures that the distribution of agents remains stationary over time because the model is analyzed in steady state. This condition implies that, across all three dimensions: savings, employment status, and age, the sum of inflow and outflow of agents at each discrete point on the distribution grid must be equal.

$$\partial_t g_t(a,e,l) = \partial_a[s(a,e,l)g_t(a,e,l)] - g_t(a,e,l)\sum_{e'\in E}\lambda_{ee'} + \\ \sum_{x\in E}\lambda_{xe}g_t(a,e,l) - g_t(a,e,l)\sum_{l'\in L}\lambda_{ll'} + \sum_{x\in L}\lambda_{xl}g_t(a,e,l) \tag{21}$$

The government's budget must balance, as expressed in eq.(22). Government revenues are generated entirely through wage and dividend taxation eq. (23), while expenditures consist of transfer expenses eq.(24).

$$gov\_surplus = 0 = revIncome - expTransf \tag{22}$$

The income tax base includes labor income, unemployment benefits, and dividends. Notably, interest income is excluded from the tax base. This is intentional and ensures that tax policy does not distort the real interest rate. During periods of negative productivity shocks, tax rates would be elevated to balance the budget. This would increase the real interest rate because the model would account for the tax rate and require a higher interest rate to offset higher taxes. Having interest income taxed potentially biases the model's core interest rate results.

$$revIncome = \sum_e \int_l \int_a \tau_1(w(z_l h_{e,t}) + \pi_l) \tag{23}$$

The transfer expense is the sum of all unemployed individuals times the unemployment/social security payment.

$$expTransf = \sum_e \int_l \int_a u_e(1-e) \tag{24}$$

The aggregate level of assets in the economy must sum to zero in the steady state equilibrium of the Huggett model eq.(25). This assumes a closed economy assumption, where every asset held by one agent corresponds to a liability held by another. In other words, for every loan extended, there must be a corresponding borrower. The real interest rate adjusts to ensure the asset market clears.

$$A(r) = 0 = \sum_e \int_l \int_a ag(a,e,l;r)da \tag{25}$$

5.3 Typical Huggett Model Setup

A standard Huggett Model will be used to compare with the heterogeneous-agent model specified above. The Hugget model will be referred to as the Hugget model, and the Heterogeneous agent overlapping generations model specified above will be the HA model.

The Huggett model has the same CRRA utility curve eq. (12). The Huggett model differs in that it has endogenous labor supply, allowing agents to choose their own hours, h, based on their asset level and idiosyncratic productivity. Labor disutility is added in eq. (26) so that the agents can pick their hours. The scale of labor disutility is denoted as $\chi . The Frisch Elasticity is 1 \frac{}{\phi_h}$.

$$d = \chi \frac{h^{(1+1/\phi_h))}}{1+1/\phi_h}; \tag{26}$$

The budget constraint eq. (27) is simpler than the HA model, as it does not require a government equation. This is because there is no unemployment, so all agents have income to make interest payments. It is important to note that the lowest ideosyncratic productivity level, z, can fully cover interest payments on the full amount of debt, leaving money over. This makes the consumption equal income from working + interest income, less savings.

$$c(a,z) = zh + ra - s(a,z) \tag{27}$$

The value function eq.(28) for the steady state is now based on the utility of the current consumption less current work, plus the value of the next asset grid level times the saving rate, plus the transition probability of moving to a new position times the difference in value between the current and future position.

$$\rho V(a,z) = \underset{c,h}{max} \ u(c) - d(h) + v'_a(a)\dot{a} + \sum_{z' \neq z} \lambda_{zz'}(v(a,z') - v(a,z)) \tag{28}$$

Decision on FOC and velocity equation The FOC is now expanded to include the marginal disutility of work so that the marginal utility of consumption equals the marginal value of saving/disaving and the marginal disutility of labor. This is seen in eq. (29). From here you can use eq. (**??**) to remove the c and h with $u'^{-1}(v'_a(a,z))$ and $(d'^{-1}(v'_a(a,z)$. This forms a viscosity solution eq. (30) that allows the value function to be computed.

$$u'(c) = v'_a(a,z) = d'(h) \tag{29}$$

$$\rho v(a,z) = \max_{c} \; u(u^{'-1}(v_a^{'}(a,z))) - d(d^{'-1}(v_a^{'}(a,z))+$$
$$v_a^{'}(a)(zd^{'-1}(v_a^{'}(a,z)) + ra - u^{'-1}(v_a^{'}(a,z)))+$$
$$\sum_{e' \neq e} \lambda_{ee'}(v(a,z') - v(a,z')) \tag{30}$$

The borrowing limit was set to match the household debt level in the HA model. This enables comparisons between the models. It is important to note that agents with the lowest level of ideosyncratic productivity have enough income to pay their interest and have income left over to consume.

This model is now much simpler as we have a simpler KFE eq. (31) equation as well as just the asset market clearing condition equation eq. (25).

$$\partial_t g_t(a,z) = \partial_a[s(a,z)g_t(a,z)] - g_t(a,z)\sum_{z' \in Z} \lambda_{zz'}+$$
$$\sum_{x \in Z} \lambda_{xz} g_t(a,z) \tag{31}$$

5.4 Saving Motives

Each agent in the model faces three primary savings motives. First, there is the intertemporal savings motive, which depends on the relationship between the real interest rate r and the rate of time preference ρ. When $r < \rho$, agents are incentivized to borrow, as saving for future consumption is relatively less valuable. Conversely, when $r > \rho$, agents are incentivized to save as the return on the asset is greater than their rate of time preference. In the case where $r = \rho$, agents are indifferent between saving and borrowing. In this model, the steady-state equilibrium satisfies $r < \rho$, encouraging agents to borrow more and save less. Second, agents save to smooth consumption over the life cycle, due to the concavity of the utility function. This leads to the typical life cycle pattern of borrowing when young, saving during middle age for retirement, and dissaving in retirement. Third, agents engage in precautionary saving behavior in response to income uncertainty, which in this model arises from unemployment risk. This motive arises because the model employs a Constant Relative Risk Aversion (CRRA) utility function, which implies convex marginal utility $u''' > 0$. Uncertainty related to future employment status drives this behavior, particularly over the extensive margin. As unemployment increases, the probability and duration of unemployment also rise. The reduction in total hours worked along the extensive margin lowers the value function not only for unemployed agents, but also for employed individuals with limited asset holdings, who face heightened unemployment risk and lack sufficient savings to self-insure against periods of unemployment. As a result, agents exhibit a stronger precautionary savings motive, aiming to build positive asset positions from which they can draw down resources during future unemployment spells, thereby avoiding borrowing constraints. As a result, the precautionary motive becomes a key driver of the differences observed between the two labor adjustment margins.

It is important to point out that while the Huggett Model has lower idiosyncratic productivity levels that can mimic unemployment's low income, and the precautionary motive that goes along with it. The distribution of ideosyncratic productivity doesn't change, and this lacks the increased precautionary motive that will be felt in aggregate when total factor productivity is down, and more agents are unemployed. In summary, my model will have precautionary savings motive vary when unemployment does.

5.5 Steady State

Figure (9) presents the steady-state consumption and saving functions for four representative groups: employed 20-year-olds, employed 45-year-olds, retired 70-year-olds, and unemployed 20-year-olds. These functions are further analyzed using a mesh diagram that depicts savings behavior as a function of all of the age and asset grid levels for both the employed and unemployed. The mean asset holdings by age closely replicate the empirical profile of asset accumulation over the life cycle. Similarly, the model's unemployment rates are calibrated to match observed patterns. For instance, the unemployment rate for 20–24-year-olds is approximately 1.5 times the overall unemployment rate, consistent with BLS data. This calibrates the unemployment rate assigned to 20-year-olds in the model. The 25–34 age group is modeled to reflect the overall unemployment rate, with both 25- and 30-year-olds assigned values aligned with the aggregate data. For all remaining age groups, unemployment rates are relatively constant and below the overall unemployment rate.

The model also produces a parabolic before-tax income profile by age, making all agents within age and employment status homogeneous, with all individuals aged 65 and older classified as unemployed/retired. Although the income gap between employed and unemployed 20-year-olds is not substantial in absolute terms, the value of employment is significantly higher. This stems from the fact that employed individuals are permitted to borrow against future income, which enables them to increase current consumption and enhances their overall utility.

The asset distribution curve displays a large mass of individuals clustered around zero assets. This concentration is driven by both precautionary saving and life-cycle consumption-smoothing motives. Young employed agents tend to borrow to smooth consumption across their lifetime. Two distinct borrowing kinks emerge in the distribution: one near -3, where the 20-year-olds become borrowing constrained, and another near -5, where 30-year-olds hit the final grid constraint of -5. It is important to note that the majority of 20, 25 and 30-year-olds do not hit these borrowing limits, indicating that only a subset of employed young agents are severely liquidity constrained.

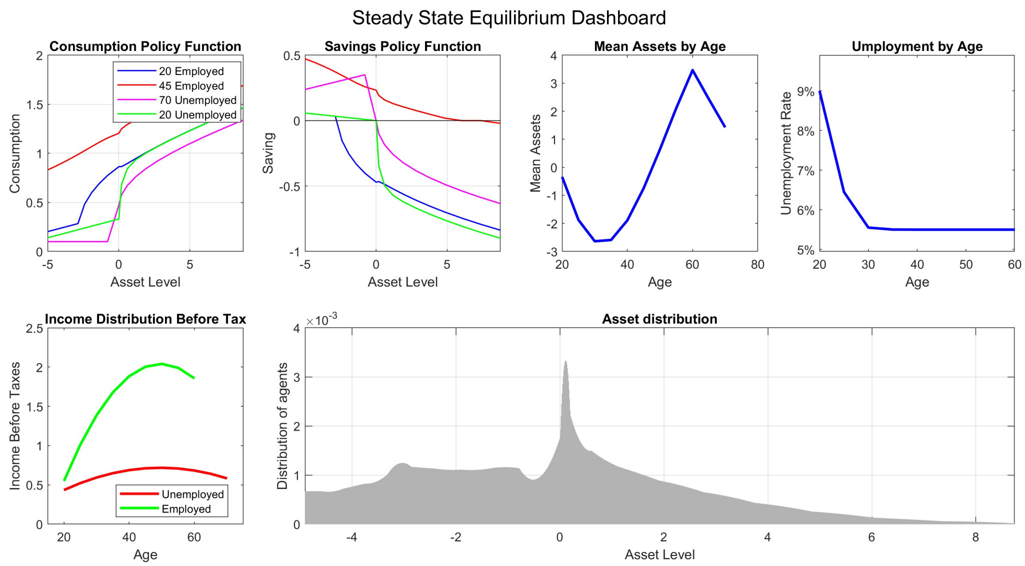

Figure 9: Steady State Dashboard

Figure (10) shows the dashboard for the Huggett model used for comparison against the HA model. The model is calibrated to have the same steady state debt levels and interest rate. To accomplish this, the minimum asset grid is moved up and down to increase or decrease the household debt. The rate of time preference is also adjusted to match the real interest rate. The main difference here is that there is now an hours policy function that determines how much each agent works, and the income distribution is normal. The consumption policy function and savings policy function show that households with higher idiosyncratic productivity tend to spend and save more than households with lower idiosyncratic productivity. Households with higher wealth levels tend to spend more and save less than households with less wealth at the same productivity level. A big departure from this paper's HA model is that agents pick their own hours. Agents with higher idiosyncratic productivity work more to take advantage of the temporary productivity boost, while wealthier agents work less, as they already have large assets to smooth future idiosyncratic productivity shocks. One thing to point out is that the hours worked of the lowest-income group increase significantly around the budget constraint as they lower consumption and increase savings, because they do not want to hit the borrowing limit and have a low income with a significant portion of income going to interest. They would rather consume less now so that the odds of hitting the debt limit will drop. The income distribution has 11 income levels to match the 11 age groups in the HA model. They are distributed normally. The asset distribution shows the missing middle as a flaw in heterogeneous-agent modeling. The distribution is like a water drop, with a higher concentration of agents with negative asset levels, and this decreases as wealth increases. In reality, the highest concentration of assets is around zero, which matches the distribution of the HA model.

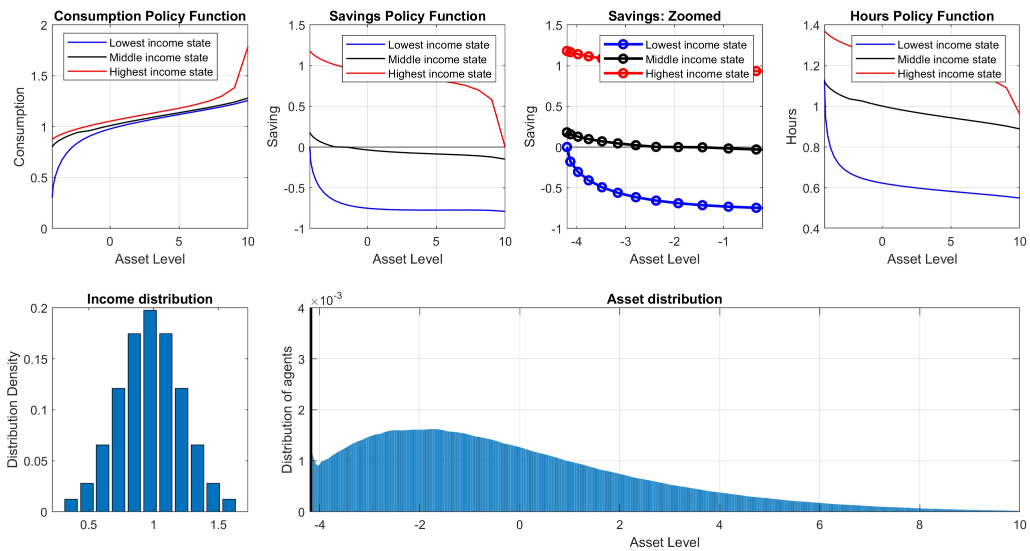

Figure 10: Huggett Steady State Dashboard

Figure (11) presents the equilibrium savings behavior for all employed agents. This does not include agents 65 or older because they are all retired. Notably, the asset axis in the figure is ordered from high to low. This display choice facilitates clearer visualization of the savings function's structure.

One key observation is that younger households tend to exhibit negative savings, reflecting borrowing. The points at which dissaving transitions from negative to zero indicate where borrowing constraints bind. A visible ridge forms around zero assets for young agents, representing agents who deliberately spend less to maintain positive asset balances for longer periods, thereby smoothing consumption in the event of future unemployment. The precautionary motive for saving drives this phenomenon. Another distinct feature occurs among agents aged approximately 35 to 60. Within this age range, individuals with lower asset holdings tend to save more aggressively. These middle-aged households face a life cycle incentive to reduce debt and accumulate savings for retirement because their labor income remains high relative to other periods of their life cycle. Additionally, some wealthy 60-year-old agents begin to dissave as their income drops before formally retiring, particularly those with substantial asset holdings. In contrast, 60-year-olds with large negative asset balances are subject to mechanically imposed higher savings rates. This modeling choice reduces the likelihood that agents retire and ultimately die with negative asset positions. This preserves the model's integrity by forcing agents to repay their debt.

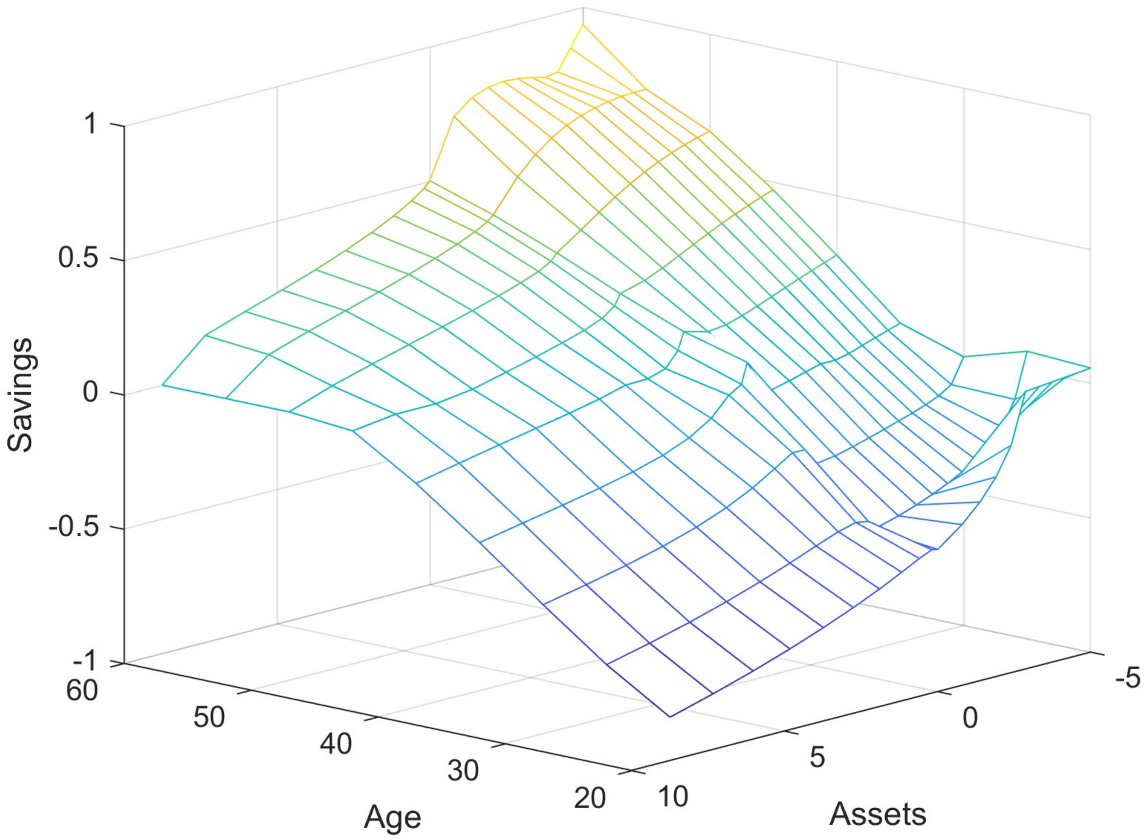

Figure 11: Employed Equm Savings Graph

In figure (12), the borrowing constraint tied to employment status is made clear by the savings behavior: once the constraint binds, savings are strictly non-negative. This reflects the model's assumption that unemployed individuals are not permitted to borrow. Unemployed agents exhibit positive savings. This is due to the amortizing nature of their debt obligations, which require not only interest payments, but principal repayment. These payments increase with the size of the loan and the agent's age, leading to higher required savings among older agents with outstanding debt. A notable pattern emerges among older agents with large negative asset balances. Agents with larger debts save less than their counterparts with smaller debt positions. This is driven by the minimum consumption constraint of 0.1. These agents consume the minimum quantity and allocate the remaining income, after covering interest payments, to savings. However, agents with greater debt levels experience lower cash flows after interest obligations, reducing their capacity to save.

In contrast, agents tend to dissave until they approach the zero-asset constraint. A spike in dissaving is observed at age 70, reflecting a $\frac{1}{20}$ probability of death and the limited bequest motive.

The bequest motive is the value of being 20 with the specific asset level. The value difference between asset grids is much smaller, as 20-year-olds can borrow to spend if they have a job. It is important to note, however, that the average 70-year-old holds approximately 0.5 units of assets. While dissaving behavior intensifies at this age, only a small proportion of agents have enough assets to spend down rapidly.

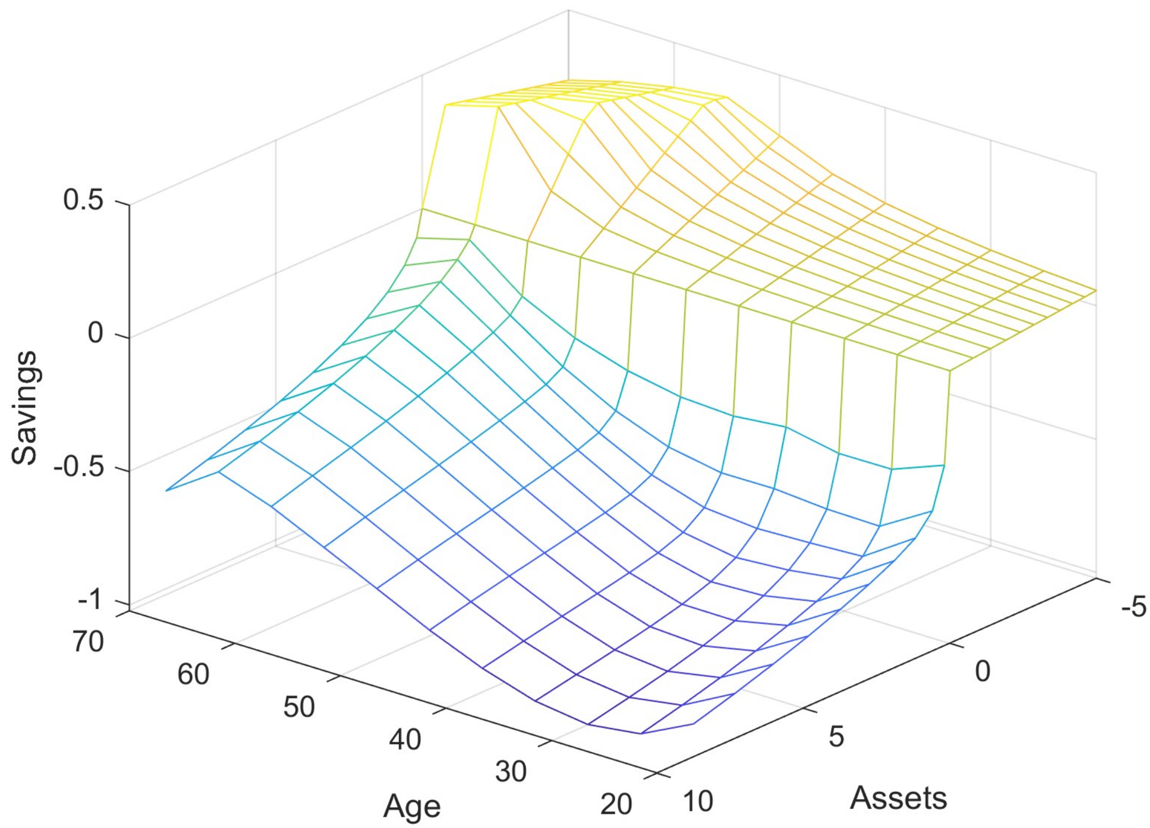

Figure 12: Unemployed Equm Savings Graph

Figure (13) shows the equilibrium savings behavior for all agents in the Huggett model. Only one graph is needed, as there are only 2 dimensions: idiosyncratic productivity and asset level. One thing to point out is that when at the highest asset level, there will be at most zero savings. Savings drops fairly smoothly with productivity level. Savings increase rather slowly as the asset level declines until the borrowing limit, after which they increase sharply across all levels. For high-productivity agents, the savings are higher, pushing them further away from the borrowing limit in the event of a negative productivity shock. For the low-income agent, savings increase to zero because they cannot borrow anymore. The HA model has many more complex shapes, based on borrowing limits and the zero-asset limit. The key difference is that you see savings for older agents increase dramatically around the zero asset line, and for young agents, the savings increase dramatically near the borrowing limit. The saving of the unemployed increases drastically around zero. Because of these differences and the movement throughout the life cycle, the missing middle is no longer missing.

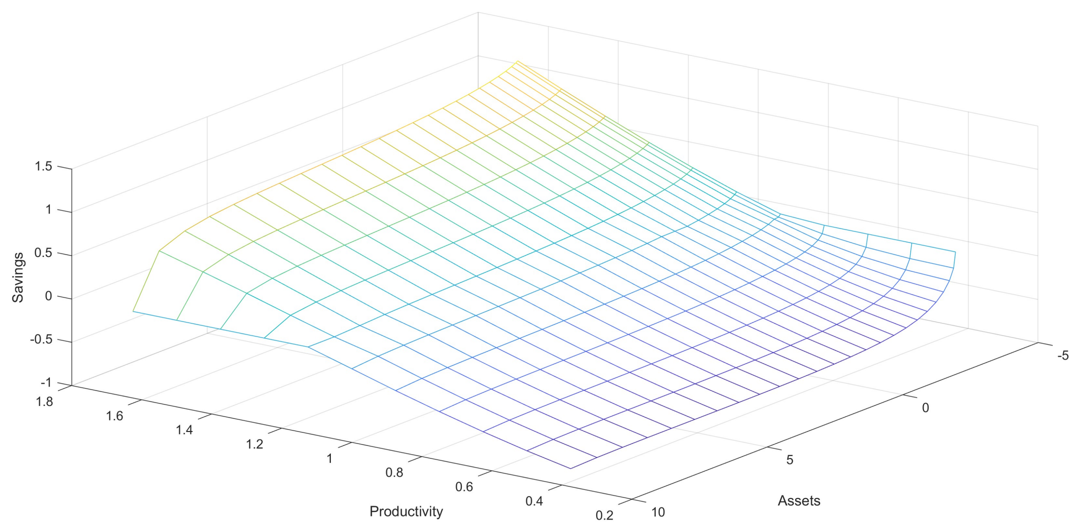

Figure 13: Huggett Equm Savings Graph

5.6 Market Shock

The effects of an aggregate productivity shock on total factor productivity A_t is analyzed over the intensive and extensive margin to identify changes to both loan demand and loan supply. Productivity shocks reduce the marginal product of labor (MPN), thereby decreasing labor demand. The impact on employment depends on whether the adjustment occurs along the intensive margin (hours worked by agents) or the extensive margin (number of employed agents). When the extensive margin is adjusted, utility levels decline due to the inability to hedge income in incomplete markets. This increases precautionary savings behavior.

In contrast, adjustment through the intensive margin leads to a less dramatic shift in loan demand and loan supply. To illustrate these dynamics, the savings policy function is compared when hours adjust (intensive margin) and when employment adjusts (extensive margin).

Figure (14) illustrates the structure of the labor market in the model. Labor supply is represented by a vertical line, reflecting the assumption that all individuals below the age of 65 want to work, while those aged 65 and above are retired. For simplicity, it is assumed that each working agent supplies one unit of labor, and there is either no disutility associated with working or the disutility of working is smaller than the utility gained from working, so that all agents under 65 will choose to work. The labor supplied is upward sloping, consistent the macroeconomic frish elasticity.

Labor demand eq.(32) is determined by the marginal product of labor (MPN), which depends on total factor productivity. In this setting, firms take the efficiency wage as given and adjust the quantity of total labor demanded in response to productivity shocks, either by reducing hours worked by agents (the intensive margin) or by reducing the number of employed agents (the extensive margin). The model assumes that labor demand always falls to the left of labor supply, resulting in an exterior solution. The intersection of the efficiency wage and MPN determines total employment, with the gap between labor supply and labor demand representing involuntary unemployment.

$$N_t = \left(\frac{(1-\alpha)A_t}{w_t} \right)^{\frac{1}{\alpha}} \tag{32}$$

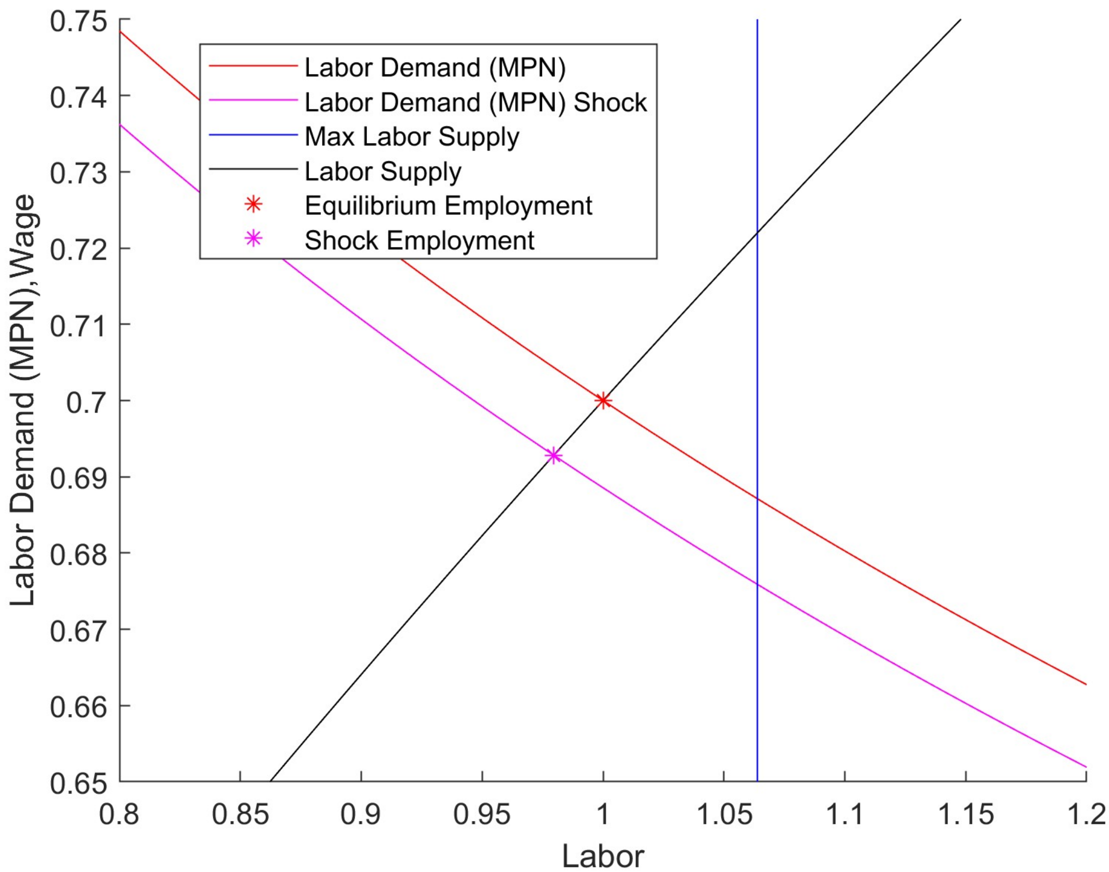

Figure 14: Labor Market

Figure (15) illustrates how utility responds to labor market shocks in an incomplete markets environment. This simplified setting assumes agents earn an income of 1 when employed (working one unit of labor) and 0.4 when unemployed. The agent consumes their entire income. This hand-to-mouth behavior abstracts from saving but will highlight the implications of labor market uncertainty.

Referencing the labor market framework in Figure (14), a negative productivity shock reduces labor demand, forcing firms to either cut hours (intensive margin) or reduce the number of employed workers (extensive margin). If all agents reduce their hours uniformly, or if all agents share the risk completely, consumption declines smoothly along the blue utility curve toward the complete market benchmark (represented by the open circle). In contrast, in an incomplete market, agents face unemployment risk without the ability to insure; utility falls to the incomplete market outcome (black star) along the red curve as the expected utility is now based on the weighted values of being employed or unemployed. This reflects heightened exposure to income volatility caused by unemployment and a strong precautionary savings motive.

While this model assumes no savings behavior, it underscores an incentive for agents to self-insure by accumulating assets. In a more realistic setting, when agents anticipate rising unemployment, their precautionary motives increase, leading to changes in saving behavior. Despite this increased precautionary motive, younger individuals continue to borrow for several reasons. First, they earn less and aim to smooth consumption over their life cycle. Second, their income is closer to the level of unemployment benefits, reducing the intensity of the precautionary motive. Third, intertemporal preferences encourage borrowing because the real interest rate is below the discount rate. Lastly, they typically have negative net wealth and must repay debt before they can accumulate precautionary savings. The reduction in borrowing among young individuals is driven primarily by the borrowing constraint faced during periods of unemployment, as highlighted in the OLG model, and by the savings behavior observed in Figure(12) of the HA model.

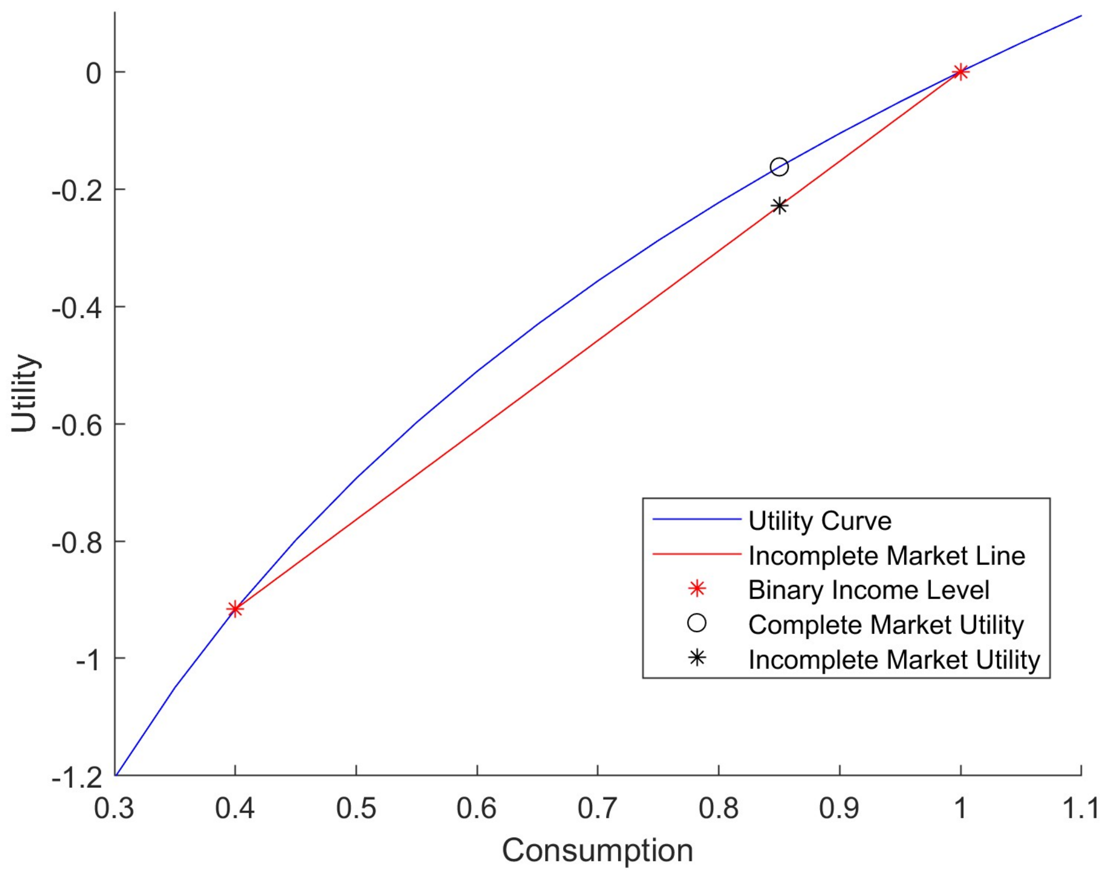

Figure 15: Incomplete Market

In most existing heterogeneous agent literature models, agents adjust their labor supply on the intensive margin. The output does not exhibit diminishing returns to labor. However, empirical evidence suggests that this is not the dominant adjustment mechanism during downturns. This motivates the modeling of unemployment along the extensive margin, where employment status changes discretely (employed or unemployed). While traditional HA models rarely explicitly incorporate dynamic unemployment risk, they are well-suited to capture the precautionary behavior observed when agents face such risks.

Figure (16) highlights a sample labor market decision that agents would face for a temporary aggregate productivity shock if agents spend all income. The horizontal line drops, and agents choose to work less. The labor supply curve is modeled for an instantaneous shock that would have the most elastic labor elasticity. As the shock's persistence increases, the change in hours decreases. If agents have higher (lower) ideosyncratic productivity, then the horizontal line is higher (lower). As agents' assets rise, the labor supply line rises, thus lowering the quantity of labor provided. This type of response would lie along the blue line in Figure (15) and would not provide the same precautionary motive as having some probability of being employed vs. unemployed. This highlights how current Huggett heterogeneous agent models fail to capture changing risk aversion dynamics.

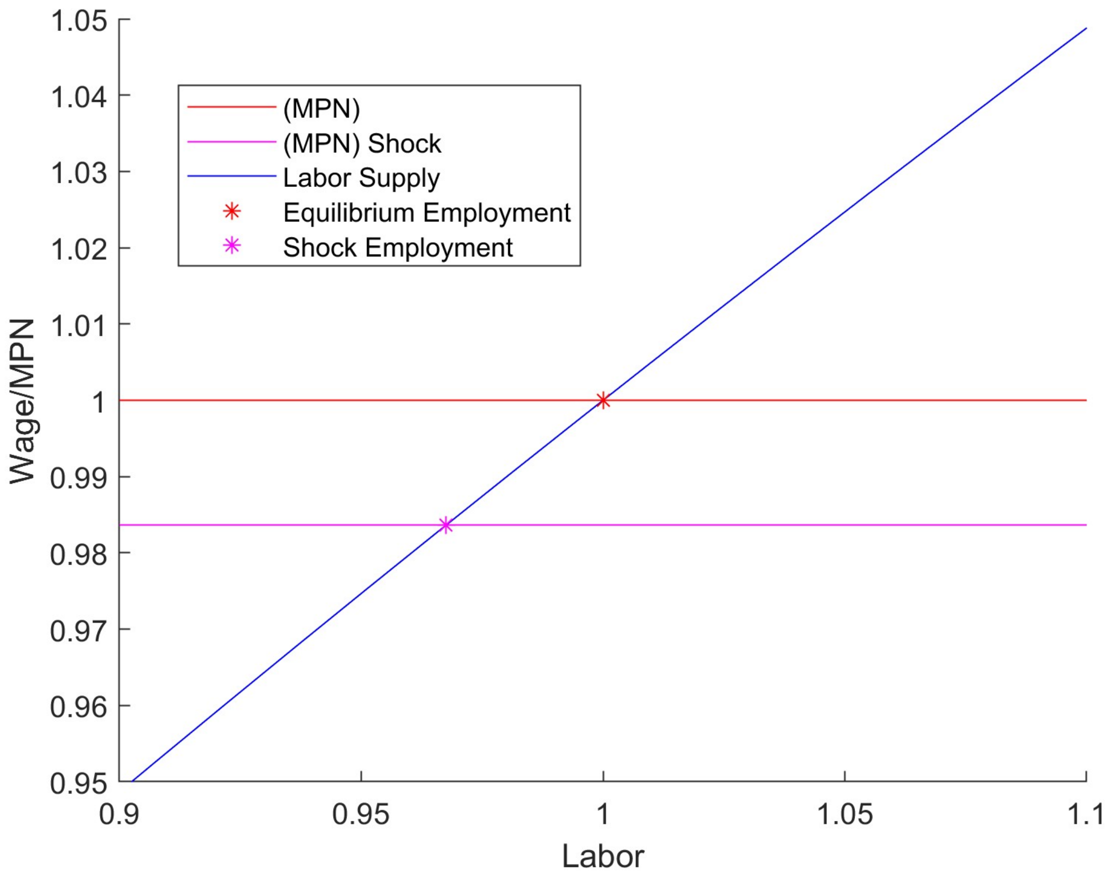

Figure 16: Huggett Labor Supply Decision

Figure (16) shows what happens when there is a permanent productivity shock, assuming agents spend all income. You can see that the vertical line productivity line drops the same amount, but now the labor supply curve shifts. The labor supply schedule eq. (33). The equation simplifies because we are assuming log preferences $\sigma = 1$ and consumption equals productivity times hours worked $C_t = A_t * N_t$. This gives us a labor supply equal to 1 $N_t = 1^{\frac{1}{1+\phi}}$.

$$w_t = \psi C_t^{\sigma} N_t^{\phi} \tag{33}$$

This feature means that if productivity changes, labor supplied changes less as the shock becomes more permanent. This is a key point of emphasis: the HA model in this paper and the Huggett model differ in that the quantity of labor can drop permanently, across both extensive and intensive margins, whereas the Huggett model does not. If HA model adjusts labor over the extensive margin, this allows for the percentage of unemployed agents to change and captures a change in the precautionary savings motive while in the steady state.

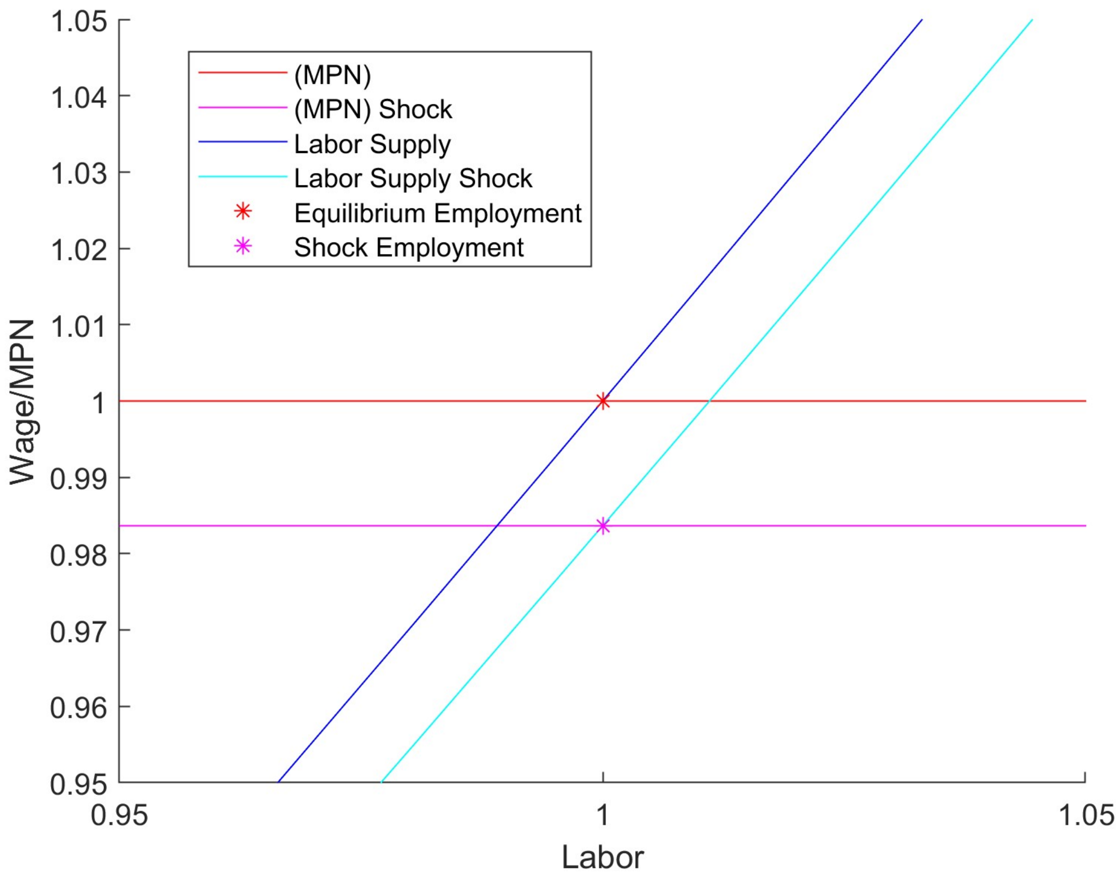

Figure 17: Huggett Labor Supply Decision in SS

5.7 Intensive vs Extensive Steady State

In this section, comparative statics will be conducted using the steady-state intensive margin with lower productivity and the extensive margin with lower productivity. This is used as the first way to motivate the impact of differences caused by the extensive margin. Adjustments to both the intrinsic (hours worked by agent) and extensive margins (employment adjustment) begin with the same steady-state calibration, incorporating the same base unemployment by age, the efficiency wage, age-specific productivity, the employment flow methodology, effective labor, the profit function, borrowing constraints, and market-clearing conditions. This provides a realistic benchmark from which to analyze how adjustments along the intensive and extensive margins impact macroeconomic outcomes from a realistic steady state. The objective is to isolate and examine the distinct effects of each margin on loan supply, loan demand, and the broader economic implications.

Both margins are studied as deviations from the steady state. A negative productivity shock is introduced, reducing the marginal product of labor (MPN) and decreasing labor demand. Following the shock, two separate adjustment mechanisms are considered: one in which only hours worked are reduced (intensive margin), with all agents working less uniformly, and another in which the number of employed agents is reduced (extensive margin). The extensive margin is adjusted by simultaneously changing the hiring and subsequently the separation rates.

When the adjustment occurs over the intensive margin, hours worked decrease uniformly for all employed agents, while unemployed agents continue to supply zero hours. The adjusted labor used falls as it is a function of aggregate productivity, as seen in Eq. (34). The dividend income of agents reduces, leading to a decline in before-tax income for both employed and unemployed individuals. However, the employed experience a more substantial reduction due to a direct decrease in hours worked and, subsequently, earnings.

$$N_t = A_t^{\frac{1}{\phi+\alpha}} \tag{34}$$

Figure (18) displays the model's outcomes after the intensive margin is adjusted. Although the overall consumption and saving profiles remain broadly consistent with the steady state, there is a modest decline in the consumption policy function, and the saving function converges closer to zero. Importantly, the age-specific unemployment rate and employment dynamics remain unchanged, with only hours worked (the intensive margin) affected. In terms of asset distribution, the structure is largely preserved, but fewer young employed individuals reach their borrowing constraint due to lower borrowing needs. Similarly, individuals with positive assets tend to hold slightly less wealth compared to the steady state.

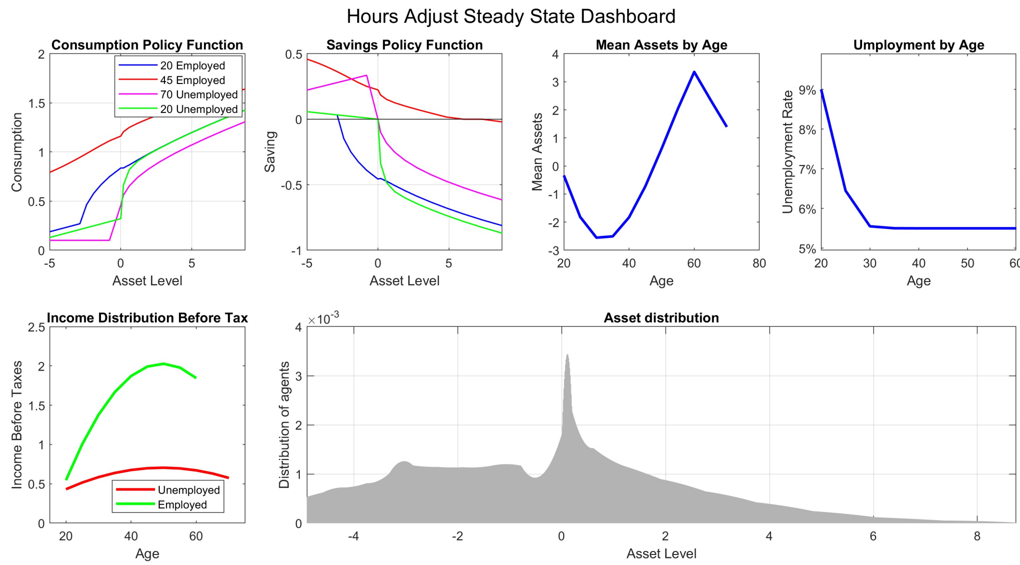

Figure 18: Intensive Adjust Dashboard

Adjustments along the extensive margin are implemented by modifying the hiring and separation rates. These rates are interlinked such that a 3% reduction in the hiring rate corresponds with roughly a 1% increase in the separation rate, following the notion from (Shimer 2012) that 77% of employment fluctuations are caused by changes in the hiring rate. The elevated separation rates for 20 and 25-year-olds remain unchanged, as does the proportion of retiring agents whose positions are filled by younger workers. The hiring and separation rate parameters are calibrated to ensure that effective labor aligns with the firm's desired labor demand.

Figure (19) displays the model's outcomes after the intensive margin is adjusted. It is important to note that not only do separation probabilities increase, but the duration of unemployment also rises, driven by both reduced hiring activity and a larger pool of unemployed agents waiting to get hired. As a result, the overall unemployment rate increases, particularly among younger cohorts, with 20 and 25-year-olds disproportionately affected. While pretax wage income remains unchanged for both employed and unemployed individuals, dividend payouts decline as firm profits decline. The consumption and saving policy functions experience mild shifts, including a noticeable kink near zero assets for employed 20 and 45-year-olds, reflecting the intensified precautionary savings motive. In response to heightened unemployment risk, those with little to no assets have a stronger precautionary motive to save.

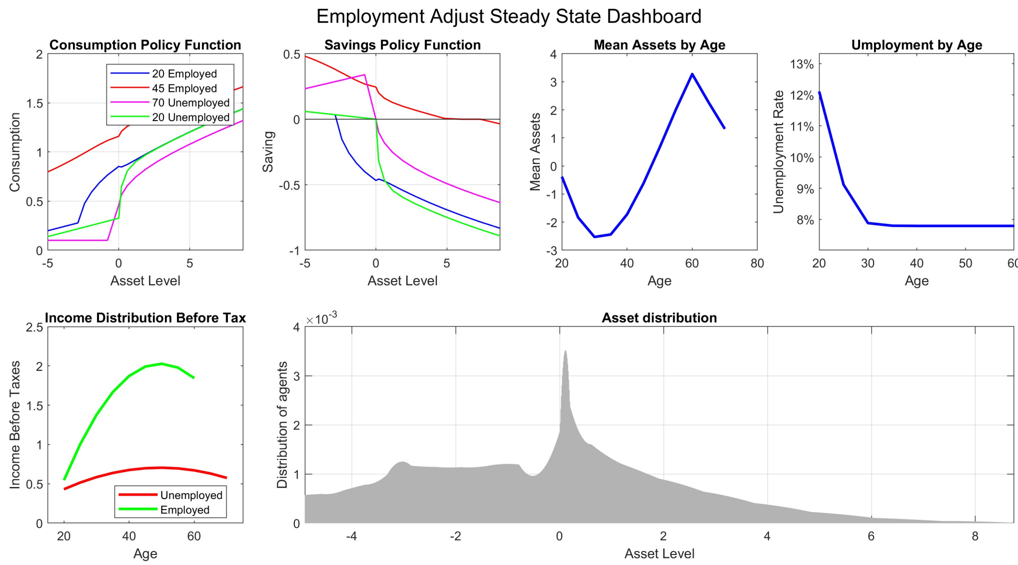

Figure 19: Extensive Adjust Dashboard

Figure (20) illustrates the shifts in loan demand and loan supply in response to a permanent aggregate productivity shock. The initial loan supply and demand curves are represented by the black lines, with the black star indicating the initial equilibrium. Loan quantities are expressed relative to the real GDP in the steady state. The real interest rate is initially below the rate of time preference, ρ, due to the strength of the precautionary savings and retirement planning motives outweighing young agents' incentives to borrow and smooth consumption. It is important to note that the loan amount is roughly 1.1 times annual GDP, 1990s, or 110% of annual GDP. This is significantly higher than the roughly 60 % to 100 % range that the US has been in since the the 1990's. This is because the model has other forms of debt, and government expenditures equal zero. First, government debt is set to zero for simplicity, and there is no corporate debt in this model. Additionally, if we set government expenditures to a non-zero number, this will lower after tax income throughout the life cycle as taxes increase to pay for spending. This will reduce the extent of income smoothing and lower borrowing from the young and saving from the middle class. While this will change the loan amount, it does not affect the direction of the shifts in the loan supplied. Because our Huggett Model does not include government spending or debt, the HA comparison will set both to zero, so we can get a better direct comparison of the baseline model rather than the effect of adding government debt and spending. An example of a loan amount graph, Figure (37), is placed in the appendix to show that there is no major change.

It is important to note that the borrowing limits by age group are fixed at the steady-state levels. This is important because the asset grid is discrete, so the debt limit would not move until the interest rate changes just enough to move the age-related debt limit. This presents a problem for two reasons. First, it makes the graph rougher as when the interest rate changes, the debt limit and hence the loan amount jump, making the line jagged. While this does not cause any problems for the theory, it does look ugly. Second, it can make solving the model more challenging. This occurs because the model is solved using quasi-Newton methods. This means that the model will expect a small change in the loan amount, but it will, in fact, be a large change when the interest rate change moves the debt limit. This leads to overshooting more iterations to computationally solve the model. The borrowing limits by age are fixed when solving the IRF, so that complexity does not go into the IRF.

The red line shows the shift in loan demand following a -1% productivity shock, with firms responding by reducing hours worked across all employed agents (an intensive margin adjustment). This results in a modest leftward shift in the demand curve. With reduced productivity, agents have a lower income across the life cycle, diminishing the need for intertemporal consumption smoothing for the young agents.

The dark blue line shows a decrease in loan supply under the same productivity shock and intensive margin adjustment. Here, the loan supply shifts slightly to the left due to two primary effects. First, the reduction in income leads to a decrease in the quantity of lending by middle-aged agents, as they have less income available to save for retirement consumption. The second reason is that there is a smaller incentive for precautionary savings. With unchanged transfers and lower employment income, the difference in the value of becoming unemployed compared to the value of being employed diminishes. This weakens the precautionary motive. However, one offsetting effect is that agents accumulate less debt. This additional wealth creates a motive to smooth some of this extra consumption potential forward.

The pink star marks the new equilibrium when the productivity shock is absorbed through the extensive margin, with firms maintaining hours but reducing the number of employed workers. This scenario results in a larger contraction in the total quantity of loans and in the real interest rate than in the intensive margin case. The model has a similar intertemporal consumption smoothing effect on the loan supply and loan demand, but now has two additional factors to take into account. First, there is a sharper decline in loan demand when aggregate hours adjust over the extensive margin rather than the intensive margin. The rise in youth unemployment primarily drives the decline in loan demand, as unemployed agents are unable to borrow. Upon transitioning to employment, these agents do not accelerate their borrowing sufficiently to offset the borrowing shortfall incurred during their period of unemployment. Second, loan supply increases when total hours are adjusted on the extensive margin rather than contracting when the intensive margin is used to adjust total hours. The heightened desire to precautionary save is caused by increased unemployment risk (extensive margin). Standard models often overlook this outcome, typically treating hours worked as the primary margin of risk, even though changes in unemployment risk are empirically more common and substantially more significant.

These findings underscore a key insight: even modest declines in productivity can generate substantial shifts in the loanable funds market, thereby affecting loan quantities and real interest rates. The adjustment of employment between the intensive margin (hours worked) and the extensive margin (employment levels) has significant implications, both empirically and within this model.

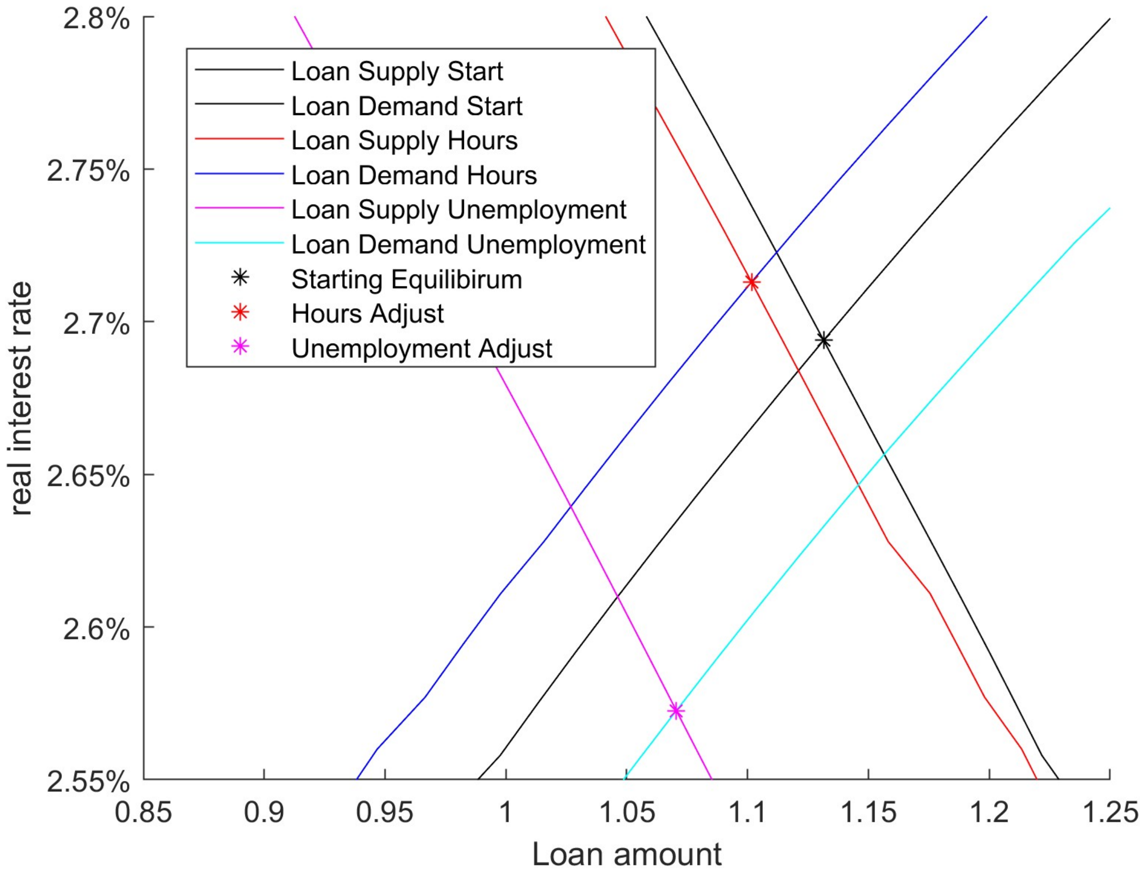

Figure 20: Heterogeneous Model Loan Demand and Loan Supply

Figure (21) illustrates the differences in saving behavior between the intensive and extensive margins for employed individuals. One of the most notable patterns is that indebted individuals aged 20 and 25 exhibit greater dissaving near their borrowing constraints under the extensive margin scenario than under the intensive margin scenario. This is due to a lower real interest rate, which means the income left over after interest payments is larger on the extensive margin, and the agent will have higher consumption and utility if they hit the borrowing limit. This means there is a diminished motive to avoid hitting the debt limit.

Between the ages of 25 and 45, the largest divergence in saving behavior occurs around zero assets. Younger individuals, who anticipate drawing down their savings, do so at a slower pace under the extensive margin. This increases the likelihood of maintaining positive asset balances in the event of unemployment. In contrast, individuals aged 30 and older with negative asset levels have a higher savings rate, enabling them to move beyond the zero-asset threshold and smooth consumption if they become unemployed. At higher asset levels, around 3, individuals with substantial positive savings tend to save less under the extensive margin. These agents are effectively self-insured against unemployment risk. This reduces the strength of the precautionary saving motive. Furthermore, in the extensive margin scenario with the lower real interest rate, the intertemporal trade-off begins to dominate the effects of the precautionary savings motive. This makes current consumption relatively more attractive to these agents, thereby reducing the incentive to save further.

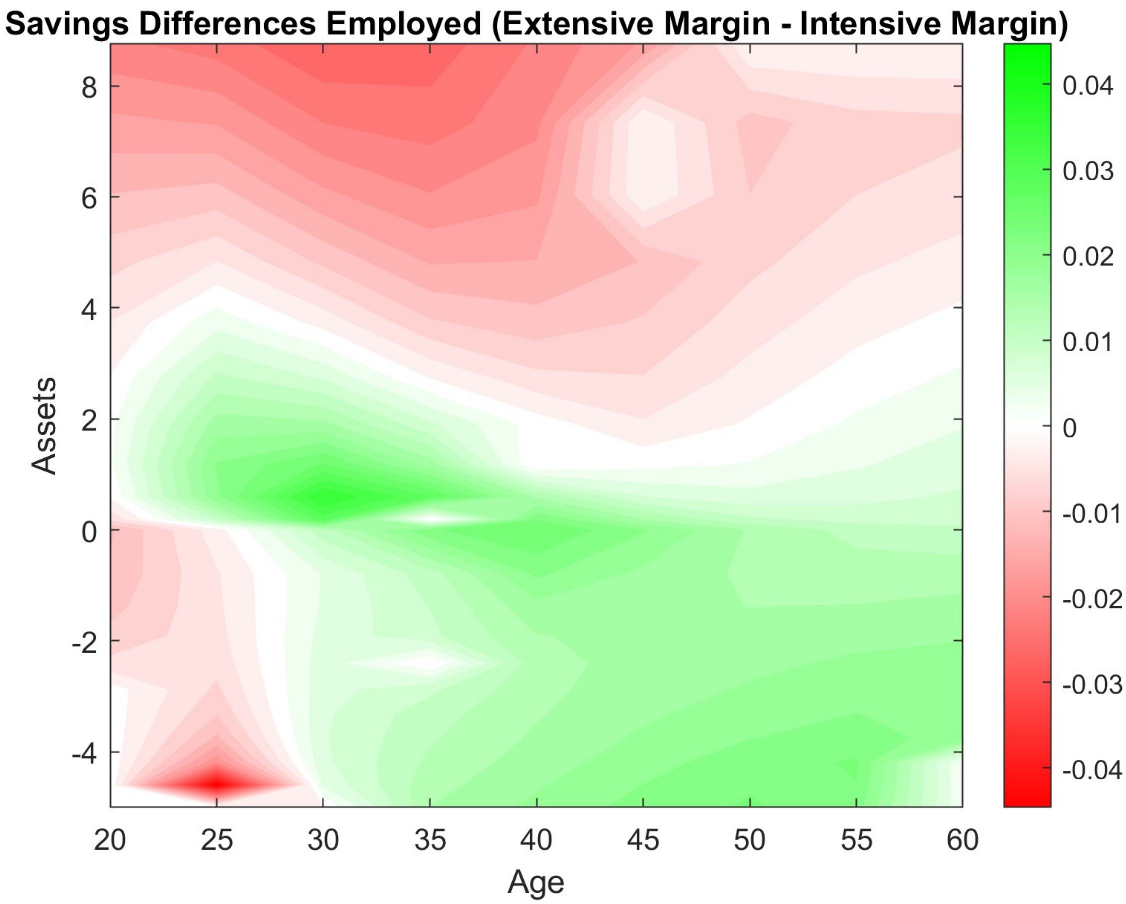

Figure 21: Savings Curve Difference Extensive vs Intensive Employed

Figure (22) illustrates the differences in saving behavior among unemployed individuals under the intensive and extensive margins. Two key insights emerge from the figure. First, unemployed individuals with sufficient assets, roughly above 1, dissave more rapidly on the extensive margin.

This is driven by a lower real interest rate, which increases the intertemporal motive to dissave as the agents benefit more from current consumption. Second, for working-age individuals with positive but low asset holdings, the saving differential is notably higher under the extensive margin. This reflects a stronger precautionary savings motive due to the higher risk of unemployment and longer average unemployment durations; agents reduce consumption to preserve resources for an extended jobless spell.

Additionally, from roughly -5 in assets for 50-year-olds to zero assets for 70-year-olds, saving rates are higher under the extensive margin. In this range, agents are constrained by debt repayment and face the minimum consumption threshold of 0.1. Because the extensive margin features lower interest rates, agents retain more disposable income after lower interest payments. This allows them to allocate more toward saving and debt reduction.

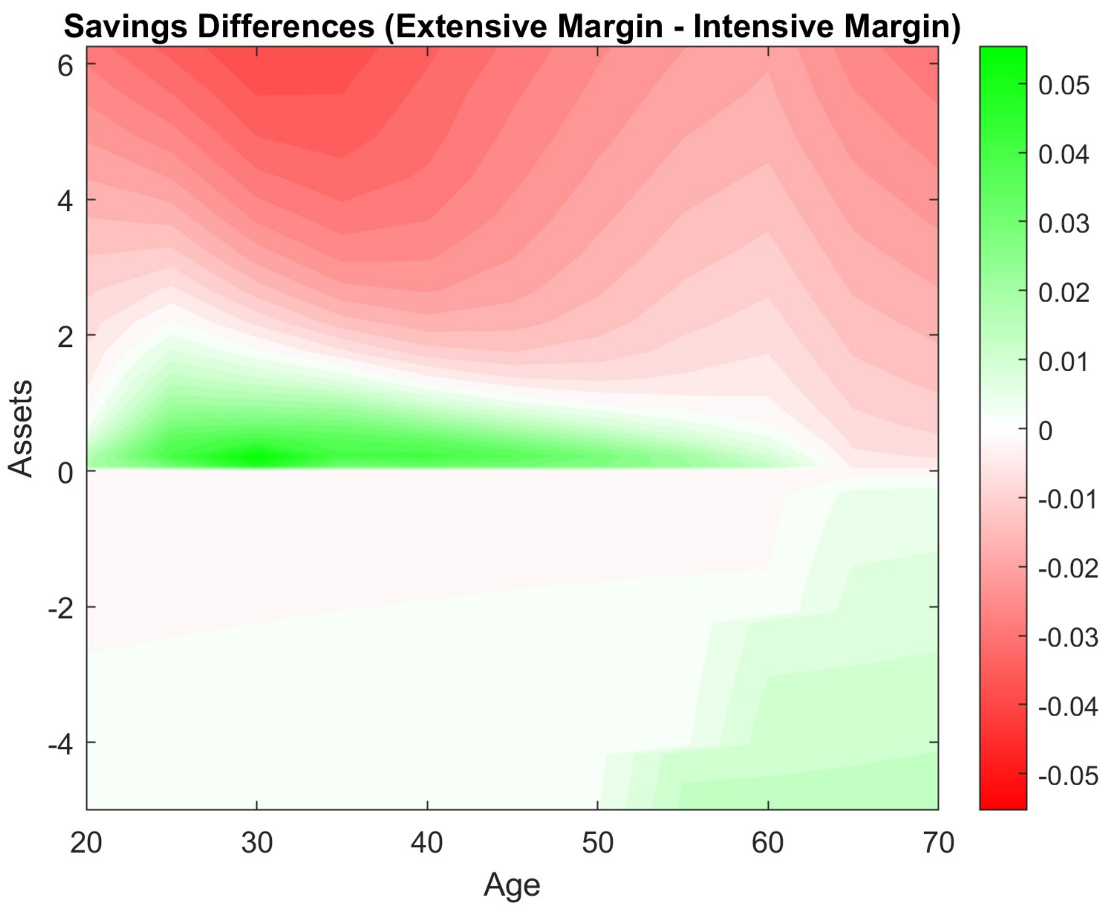

Figure 22: Savings Curve Difference Extensive vs Intensive Unemployed

5.8 Huggett model Response

The graphs of loan demand and loan supply for the HA model show significant differences in interest rates and loan quantity, driven by the extensive margin. While this was nice to see, it would make sense to see what a standard Huggett model would do so that we can see how the extensive model performs differently from the existing literature. Figure. (23) highlights the difference in the Huggett model before and after a permanent productivity shock. The size of the graph is the same as Figure (20) to show the differences on the same scale. Two things should stand out. First, the lines are less steep, so loan demand and loan supply are much more elastic. This makes sense because the standard heterogeneous agent model does not have the life cycle component that creates strong and persistent loan demand for young agents who have a lower income for years, and they do not have the incentive for the middle aged to save for retirement to smooth consumption regardless of the interest rate. The second is that the Figure (17) shows that there will be very little change in the labor provided by each household for a permanent shock. This effectively means that only the real income drops because of the productivity shock, and roughly the same hours are worked. It is important to note that hours worked increase slightly for low-income households around the borrowing limit. The only effect is that the borrowing limit with respect to income effectively goes up. This results in such a small change in the loan supply and loan demand curves that the difference is hard to see in the graph. It becomes clear that the Huggett model cannot cause dramatic changes in the interest rate in comparative statics. While not compared to capital and its effect, an Aiyagari model would have almost no difference in interest rates in response to a permanent productivity shock, and we are looking at the steady state. The interest rate response in the IRF changes when the interest rate changes due to the Euler equation, as explained further in that section.

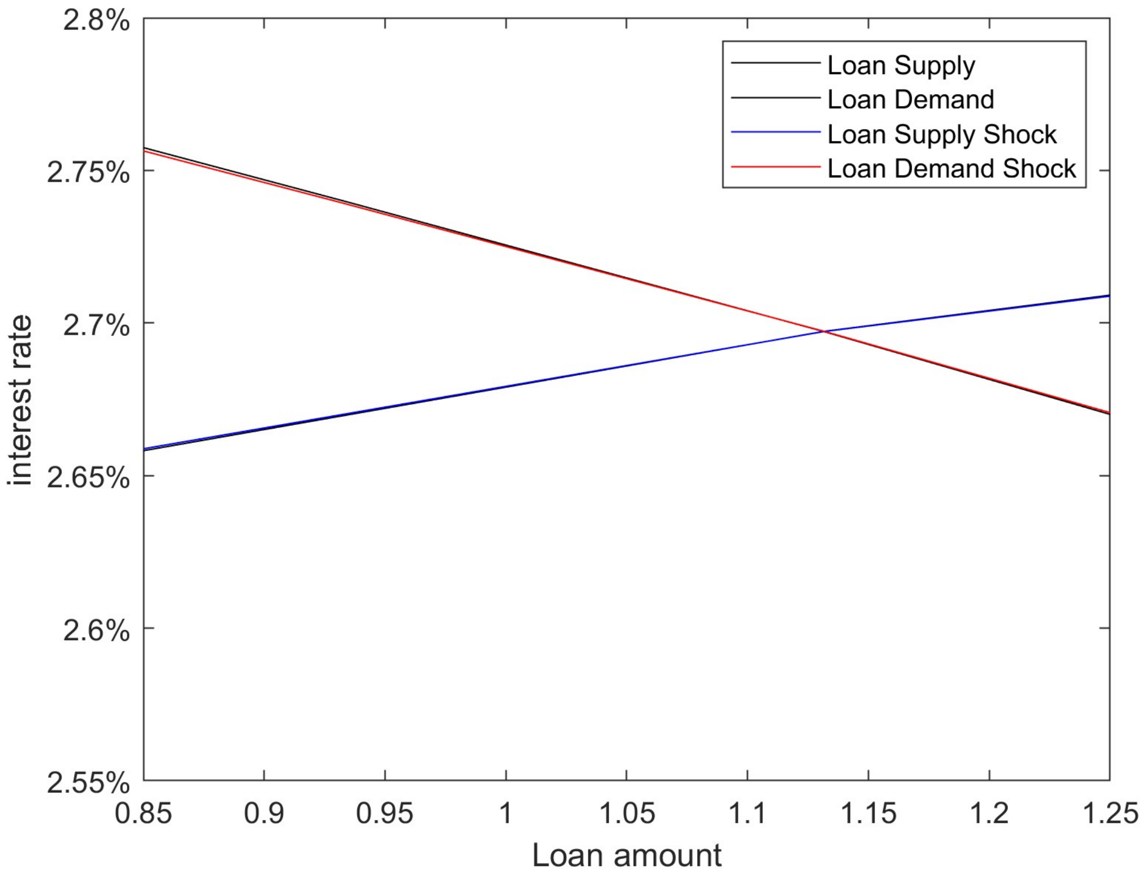

Figure 23: Huggett Model Loan Demand and Loan Supply

6 Analyzing Dynamic Responses: The impact of productivity shocks in HA framework

6.1 Overview

Comparative statics gives us a good look at what happens around the steady state, but studying the impulse response would give us a better idea of what distinguishes the modified heterogeneous agent model from a standard Huggett model. This will highlight how incorporating employment risk significantly impacts interest rate dynamics. The variables of interest are real interest rates, household loans outstanding, the percentage of individuals borrowing constrained, hiring transition probabilities, and the unemployment rate. Real interest rates and household loans outstanding will highlight developments in the loanable funds market. The hiring transition probabilities and the unemployment rate will give us a good idea of what is going on in the labor market and the extensive margin. The percentage of households constrained shows us how many households cannot borrow, and this is impacted by reaching the borrowing limit as well as being unemployed with non positive assets. This clearly shows how unemployment drives borrowing in the extensive-margin model.

Three models will be compared. The extensive and intensive model will be compared with the standard Huggett model. This will not only show the difference in the models, but also the necessity of the extensive margin.

6.2 Computation

This paper will use computational methods from (Achdou et al. 2022). They solve the HJB value function backward from the steady state to the present, then solve the KFE or the distribution of agents forward in a continuous time framework. Continuous time models can often make intuition tricky because they require you to think more closely about transition dynamics and model features. A few assumptions are made for simplification. First, we are assuming that firms are single handedly choosing the labor used based on the labor supply curve being a given. Second, we are assuming that the government balances the budget in each time period. This assumption will be relaxed in the appendix and ran. There are relatively small differences in outcomes, but the overall outcomes are the same. Because assuming no deficits is simpler and yields results very similar to ours, we will use it. The labor market is a key example of looking at transition dynamics because unemployment cannot jump by unpredictable amounts instantaneously. Because of this, quadratic labor adjustment costs were added to both the intensive and extensive models to smooth labor dynamics and better match the real world. Two papers that had quadratic labor adjustment costs are (Janko 2008) and (Lechthaler and Snower 2013). They use adjustment cost values of 0.98 and 2, respectively. A quadratic labor adjustment cost is added in line with the papers, and to avoid instantaneous changes in employment.

The sequence for solving the dynamic response for the HA models is as follows. First, the firm chooses the labor level based on the productivity shock and the adjustment cost. This can be done because, the wage curve is known and the firm maximizes profit based on the rate of time preference. None of this is affected by households. This simplification yields the desired labor level in IRF by adjusting the hiring rate in the extensive margin model and reducing all employed agents' hours equally in the intensive margin model. Second, tax rates need to be adjusted. We will know the output and the number of unemployed individuals at each stage, so the tax rate can be adjusted to ensure tax inflows match tax outflows based on the no deficit assumption. This is relaxed in the appendix and yields results similar to those obtained. Third, we need to find the market clearing interest rate. We know employment dynamics and tax responses over the cycle, so we can keep those inputs constant and focus on clearing the asset market by adjusting interest rates during the IRF. It is possible to have all three variables optimized at the same time, but it isn't necessary because the tax rate and interest rate do not affect the hiring decision of the firm, and the interest rate does not affect the tax rate, as the interest payments are not part of the household and government budget equation.

6.3 Aggregate Labor Level

Eq.(35) solves the effective labor that the firm chooses throughout all time periods to maximize the firm's utility. The utility function is the integral of all profits and labor adjustment costs discounted by the rate of time preference. For simplicity, the labor adjustment cost is an implicit cost, so it does not affect profitability; it only affects the firm's decision-making and the path that effective labor takes.

$$V = \max_{N_t} \int_0^{\infty} e^{-\rho t}(A_t(N_t)^{1-\alpha} - w_t N_t - \frac{\theta}{2}\dot{N_t}^2)\,dt \tag{35}$$

It is important to note that the wage, w_t, is based on the Frisch elasticity of the macroeconomy. Because of this, the wage curve is Eq. (36). This can be substituted into the firm's profitability equation so that the firm can maximize profit by adjusting N_t throughout the different time frames.

$$w_t = (1-\alpha) * N_t^{\phi} \tag{36}$$

Figure (24) shows the path that the effective labor takes both with and without adjustment costs. Without adjustment costs, the production shock is mirrored; with adjustment costs, the response is slower, and the optimal labor response is dampened. This also removes a large, instantaneous jump shock in effective labor, making hours and employment react more smoothly and less discontinuously. The profit function goes over shows profit over time. The profit flows to households as dividend income.

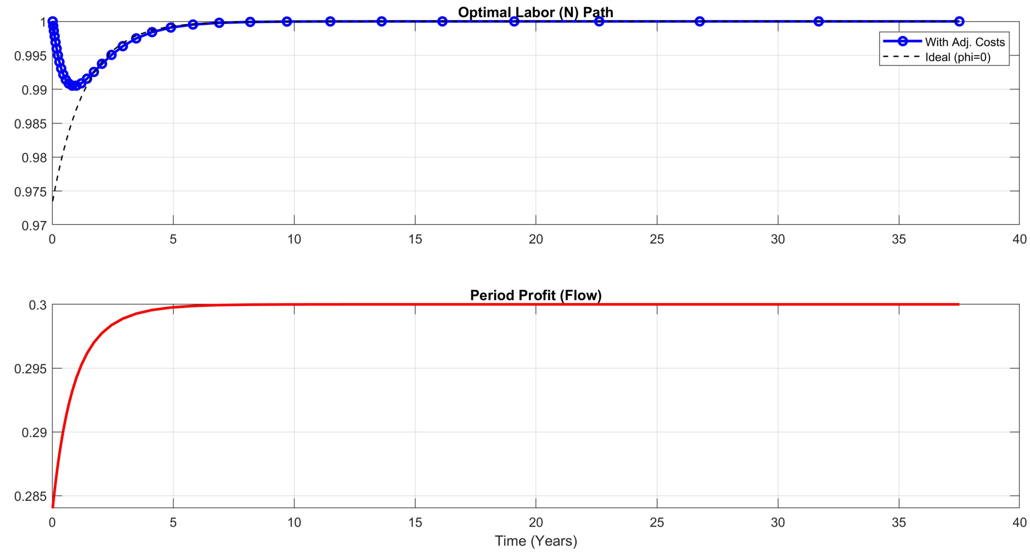

Figure 24: Labor IRF

From the optimal labor, we can adjust hours in the intensive model by setting all agents' hours equal to the optimal labor path value. No additional calculations are needed. For the extensive margin model, we will need to solve the unemployed to employed transition probability. The unemployed to employed transition probability will be adjusted for all time periods through optimization methods until the labor path in the model matches the optimal labor path. This completes the employment transition dynamics and allows us to move to taxation.

The next market clearing condition is the government budget function. For the intensive margin model, we have hours worked, profit in each time period, and the same unemployment and unemployment payments throughout the shock. This allows us to adjust the taxation as it is taxation on all income, less interest payments, which is wage income plus dividend income. This means that interest rates are not needed to solve for the tax rates. We know the unemployment and retirement expenses, so the tax rate for each time period can be solved using a Broyden update across all time periods and tax rates. In the extensive margin model, hours worked remain at 1, the same profit is earned in each time period, and we now know the changing unemployment rates and payments. From here, we can use the same Broyden method to find the tax rate.

Given the labor and tax rate dynamics, all real interest rates can be calculated. Using standard Broyden methods yielded quite jagged interest rate responses, so Gauss-Newton methods were used instead. Using Gauss-Newton methods still left a mild chop. To improve this, a few things were done. The first is that a smoothing penalty was added, and the Levenberg-Marquardt method was used to avoid overly large steps and those that would make the interest rate more jagged. The smoothing penalty is given in eq. (37). The penalty is based on the curvature of the interest rate around a central point. This means the smoothness penalty isn't applied to the first and last interest rates. Equation (38) is the smoothing Jacobian that is seen from changing the interest rate of one of the variables. It is important to note that changing the rate at a single time period not only changes the smoothness penalty for that period, but also affects the time periods directly before and after it.

$$p_i = \frac{x_{i+1} - xi}{\Delta grid_f} - \frac{(x(i) - x(i-1))}{\Delta grid_b} \tag{37}$$

$$\begin{aligned} R(i,i-1) &= \frac{1}{\Delta grid_b} \\ R(i,i) &= -\frac{1}{\Delta grid_b} - \frac{1}{\Delta grid_f}; \\ R(i,i+1) &= \frac{1}{\Delta grid_f}; \end{aligned} \tag{38}$$

From here, the Augmented F and Augmented J matrices will be updated. The F matrix is the residuals from the asset clearing equation. Equation (39) shows that the augmented F is the original F with the smoothness penalty score added below it. If the F was 5x1 and the Ps was 5x1 $F_a ug$ would be 10x1.

$$F_{aug} = [F; sqrt(\lambda_s) * Ps]; \tag{39}$$

The augmented J matrix, eq (40), is the standard Jacobian stacked on top of the penalty Jacobian. The Jacobian, J, would be 5x5, and the R would be 5x5. This means when they are stacked, J_{aug} is 10x5.

$$J_{aug} = [J; sqrt(\lambda_s) * R]; \tag{40}$$

Equation (41) represents the Hessian, which is a combination of the transpose of J_{aug} times itself and the Levenberg-Marquardt dampening parameter μ times the identity matrix. If J_{aug} is 10x5 and I was 5x5, the resulting output is 5x5, the same as the original J, but now we have damping and smoothing built in compared to a standard Hessian.

$$H = J'_{aug} * J_{aug} + \mu * I); \tag{41}$$

Equation (**??**) represents the gradient. This is calculated using the transpose of J_{aug} times f_{aug}. This gives us the dissent rate. If J_{aug} was 10x5 and F_{aug} was 10x1 this would yield a 5x1 matrix.

$$G = J'_{aug} * F_{aug}; \tag{42}$$

Equation (43) shows the change in the interest rate variables based on the Hessian and Gradient. If the H was 5x5 and the G was 5x1, this would yield a change in interest rate of 5x1, matching the number of interest rates and residuals.

$$dx = -H^{-1} * G; \tag{43}$$

The Huggett model uses the same methods to solve for the interest rates. The biggest difference between these models is that the interest rate does not directly control consumption. It is set to clear the asset market and is more of a response to intertemporal consumption and the marginal benefit of consumption at the aggregate level.

6.4 Results Comparison

From here, we will show the responses to productivity shocks with different levels of persistence on the extensive, intensive, and standard Huggett model. The different persistence levels will highlight how agents respond differently to shocks of different sizes. These graphs will show the shock size with no government debt, no government spending, and a zero budget deficit in all time periods. This gives a similar result to when we include government debt, government spending, and a government budget constraint that doesn't balance, but is much simpler and easier to understand. Figure (39) in the appendix displays the irf if government purchases of goods and services is 17% of GDP, government debt is 70% of GDP, and the government is not required to balance its budget each time period.

Figure (25) shows what happens with a productivity shock that has low persistence. First, let's look at the labor market dynamics. The unemployment rate stays the same for the intensive and Huggett model as hours are forced lower in the intensive margin, and in the Huggett model, agents actually choose to work slightly more, as the income effect dominates the substitution effect, and especially lower income agents work more hours. The unemployment transition dynamics also stand out: the intensive margin model maintains unemployment, whereas the Huggett model has none. The expensive margin model reduces the hiring rate before returning to equilibrium levels. Let's look at the borrowing constrained households. The intensive and extensive models start and end with the same level of constrained households. Constrained households are agents that cannot borrow. This includes agents that have hit the age based borrowing limit or agents that are unemployed and have zero or fewer assets. They either can't borrow because they are at their borrowing limit, or they can't get a loan because they're unemployed. The extensive margin model shows an increase in the number of constrained households as a higher percentage of agents are now unemployed and have negative asset values. It is important to note that youth unemployment will increase the most in this model, and because youth have negative asset levels, this will increase the number of constrained households. There is a mild dip in constrained agents in the intensive margin model as there is no change in unemployment, but there is a total drop in wages as hours worked drop. Agents take lower income into account and end up borrowing less. The borrowing limit does not move, as doing so would cause problems with computational methods. There is a mild increase in the constrained households in the Huggett Model as low income agents with low wealth have a higher marginal utility of consumption when their income drops and partially make up for this by increasing their borrowing and having a higher chance of hitting their borrowing limit. This differs with the HA intensive margin model because the intensive model has consumption smoothing and life cycle dependent income, whereas the Huggett model does not.

Let's look at the interest rate. It is important to note that the Intensive and Extensive Models have a quadratic adjustment cost on the rate of change of labor. This means that labor adjusts more slowly as depicted in Figure (24). The Huggett model adjusts instantly. That means the Hugget model doesn't show a negative interest rate because the decrease in earnings is instantaneous, and from that point on, income only increases, so the interest rate is above the steady-state level and then drops back to the steady state as the shock dissipates. Both the extensive and intensive margin models show that interest rates drop when the shock hits due to intertemporal substitution. From there, output bottoms out and starts increasing. Rates increase and even overshoot the steady state, as agents have perfect foresight and can see their future income rising. As the shock approaches zero, the interest rate returns to its steady-state value. It is important to note that the magnitude of the interest rate IRF is larger in the extensive model, as it exhibits stronger variation in precautionary motives.

There is also a divergence in household loans. The Huggett Model shows total loans increase before dropping. Both the Intensive and Extensive Models show that interest rates drop when agents are borrowing less, reflecting the decline in expected overall income. As agents expect their incomes to increase, households resume borrowing. This corresponds to interest rates above the steady state value. It is important to note that the extensive margin model shows a larger decrease in loans as the number of constrained households increases. When the unemployment rate goes back down, the agents don't take out loans fast enough to make up for the lack of borrowing, and household loans don't reach the same peak that the intensive margin generates. The point at which loans increase past the steady state does not align with the empirical findings later in the paper. This is because, in reality, agents do not have perfect foresight and will likely hold on to negative beliefs about the future, based on the past precautionary motive, which is stronger than what they will see in the future. This can also be seen as agents being less likely to lend and borrowing conditions being tighter. This is not included in the model.

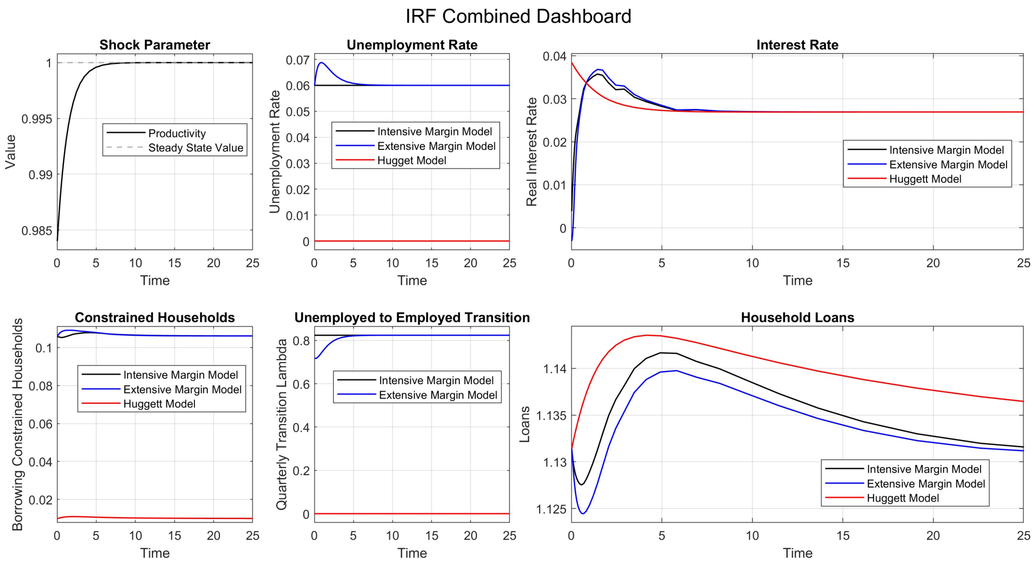

Figure 25: IRF Shock Low Persistence

Figure (26) shows what happens with a productivity shock that has a higher persistence. The first thing to notice is that the shock is the same magnitude; it just lasts longer. Now that the shock persists longer, both the Intensive and Extensive models adjust labor more quickly and for a longer period to address the more prolonged shock. This is evident in the unemployment rate going higher and taking longer to drop to steady state. Interest rates follow the same general path, but drop lower for the intensive and extensive models. When the overshoot is above the steady state, they don't jump up as high, but stay higher for longer. The constrained households and Unemployment transitions match the same theme or a larger, more persistent move. What really stands out is the Household Loans. Here, the Extensive Model not only deleverages further, but it doesn't overshoot the steady state and looks much more similar to the empirical findings later in this paper. In this model the difference between the Huggett and both HA models becomes stark as both interest rates and household loans become much different.

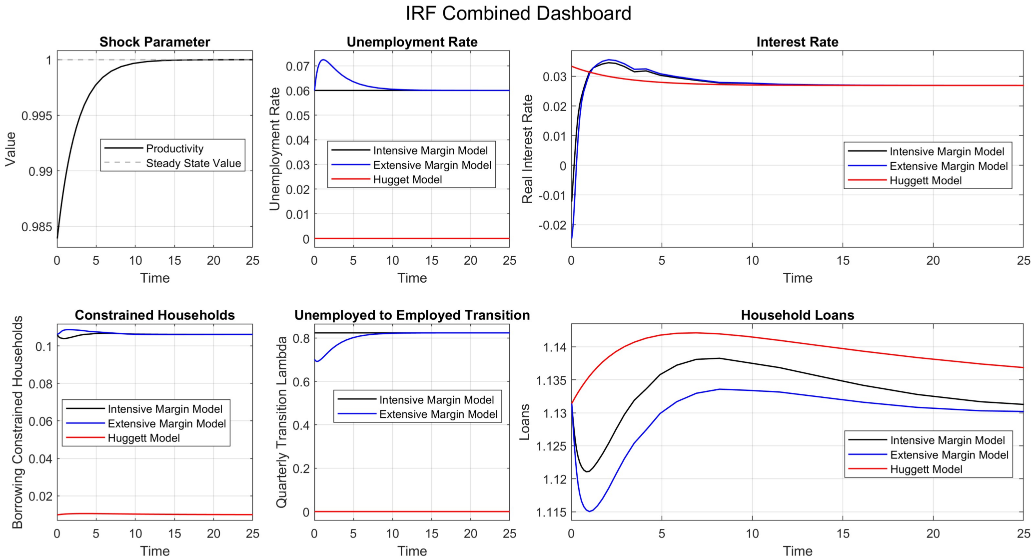

Figure 26: IRF Shock High Persistence

Figure (27) shows what happens if the market at equilibrium faces a permanent shock and the models transition to their respective steady state as seen in the comparative statics. First, the extensive margin unemployment rate increases and stays permanently elevated as the transition rate of unemployed ot employed drops. Both intensive and extensive margin models have interest rates drop before increasing to their permanent level. The intensive margin has a higher interest rate than the steady state, whereas the extensive margin has a lower value than the steady state. The Hugget Model's interest rate is roughly unchanged. Constrained households increase for the Extensive Model as unemployment goes up, go down for the Intensive margin as borrowing limits stay the same as income goes down, causing agents to borrow less, and Constrained stays roughly the same for the Huggett Model. Household deleveraging has significant differences. The Huggett Model has minimal changes to household loans and both the Intensive and Extensive Models have drops in leverage. The Extensive margin model is more pronounced. Once again, we are seeing significant differences in the two models.

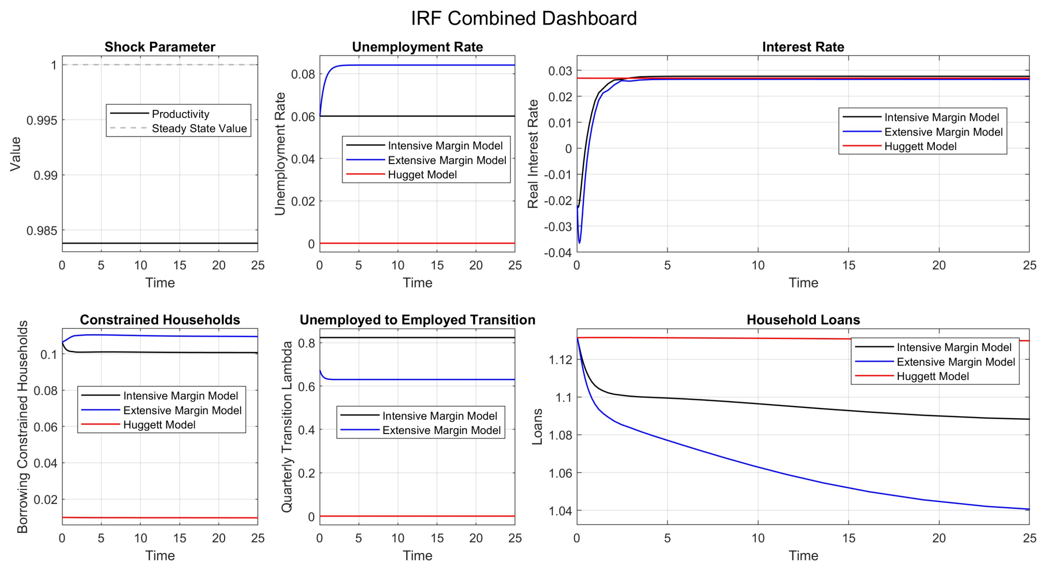

Figure 27: IRF Permanent SHock

7 Bridging Theory and Data: A VAR/LP investigation into the labor market and household debt

7.1 Overview

In this section, the dynamic responses of weekly hours worked, the unemployment rate, and real household debt based on labor productivity shocks will be analyzed. The methods used will be VAR and LP This section will be much more informative than the empirical motivation alone because it directly evaluates the channels the HA models are designed to explain. Once again, European data will be used because there is not ample state level data on US household debt and weekly hours worked. The European panel provides substantial cross-country variation that is unavailable in aggregate U.S. data. Specifically, European countries will be split into groups and analyzed for variation in weekly hours worked, unemployment, real household debt, and the household debt-to-GDP ratio. This variation provides a richer evaluation of the HA models' predictions and directly assesses the channels the HA models are designed to explain.

7.2 Data

The models use quarterly macroeconomic data from 1998 to the end of 2019, collected from several sources. This includes the earliest date across all data sources and ends just before the Covid outbreak. Labor productivity obtained from Eurostat. Labor productivity was used because multi-factor productivity was annual data, and the HA and Huggett model uses only labor, so it is more appropriate. Eurostat also provided the consumer price data, which is a harmonized measure. Average weekly hours worked also came from Eurostat and was used to distinguish fluctuations in labor input arising from changes in hours worked versus changes in employment. Unemployment data are obtained from the OECD database. Household debt data are collected from the Bank for International Settlements (BIS), measured as credit to the household sector and credit to household sector to GDP. The resulting panel combines labor productivity, weekly hours worked, the unemployment rate, and real household debt across European economies at quarterly frequency, allowing for a deeper dive into the relationship between employment dynamics and household debt amid productivity shocks in the business cycle.

7.3 Empirical Methods

The methods in this section will include detrending the data using a Hamilton Filter, followed by analysis with Vector Autoregression (VAR) and Local Projection (LP).

The filter method used was (Hamilton 2018). The Hamilton filter was used for a few reasons. First, it removes the endpoint bias because the Hamilton Filter does not include smoothing that looks forward at the starting point and smoothing that looks backward for the end point. Second, the filter is much more interpretable because the cyclical component is the deviation from the predicted path. This will allow the filter to do a better job in low frequency cycles as well as avoid serial correlation, persistence, and co-movement between unrelated variables. Lastly, there is no artificial smoothing parameter, as a Hamilton filter uses a simple regression eq. (44)

$$y_{t+h} = y_t + y_{t-1} + y_{t-2} + y_{t-3} + v_{t+h} \tag{44}$$

The Hamilton Filter traditionally looks ahead 8 quarterly periods or two years. (Jordà, Schularick, and Taylor 2017) argued for looking ahead 5 years when examining debt cycles. Because of this, both the 8-period 2-year and 20-period 5-year look-ahead will be used to cover all bases.

To ensure that the estimated dynamic relationships are not driven by a particular econometric specification, the analysis employs both VARs and LPs. The VARs follow (Sims 1980) and provide a framework for estimating impulse response functions under a structured dynamic specification. In contrast, LPs follow (Jordà 2005) estimate responses at each forecast horizon separately, making them more robust to model misspecification and better suited for exploring nonlinear and heterogeneous effects. The empirical evidence is strengthened by the fact that both methods yield economically meaningful dynamic relationships among the variables.

The IRFs are computed from combined data across all countries. The countries include: Austria, Belgium, Czechia, Denmark, Finland, France, Germany, Greece, Hungary, Ireland, Italy, Poland, Portugal, Spain, Sweden, and Norway. It is important to note that several countries were excluded from the original empirical model because they lacked household debt data, and Switzerland was specifically excluded due to very limited cyclical data from 2015 to 2019.

The VAR will be setup is eq (45). This is a VAR with 4 quarters of lag. The variable setup in eq (46) is as follows: Labor Productivity is first, followed by weekly hours worked, the unemployment rate, and real household debt. There are alternative measures of household debt that can be used and are used in the appendix. These include nominal household debt and household debt-to-GDP. Both the VAR Figure (42) and LP Figure (43) in the appendix show similar results. Figure (44) and Figure (45) show that Debt to GDP increases following the negative shock because GDP drops more quickly than the household debt. As time proceeds, the debt-to-GDP ratio drops slightly below GDP before returning to normal.

There is a built-in AR(1) process for Labor Productivity to make it easier to model the shock to the structural model. This effectively puts zeros in all the first row of Φ, excluding the a_{11}. This can be seen in eq. (47) and shows that Labor productivity is an AR(1) and all of the other variables follow a traditional lagged VAR structure.

$$(I - \Phi_1 L - \Phi_2 L^2 - \Phi_3 L^3 - \Phi_4 L^4) X_t = e_t \tag{45}$$

$$X_t = \begin{bmatrix} A_t : \text{Labor Productivity} \\ H_t : \text{Weekly Hours Worked} \\ U_t : \text{Unemployment Rate} \\ B_t : \text{Real Household Debt} \end{bmatrix} \tag{46}$$

$$\Phi_1 \text{ layout } \begin{bmatrix} a_{11} & 0 & 0 & 0 \\ a_{21} & a_{22} & a_{23} & a_{24} \\ a_{31} & a_{32} & a_{33} & a_{34} \\ a_{41} & a_{42} & a_{43} & a_{44} \end{bmatrix} \quad \Phi_{2,3,4} \text{ layout } \begin{bmatrix} 0 & 0 & 0 & 0 \\ a_{21} & a_{22} & a_{23} & a_{24} \\ a_{31} & a_{32} & a_{33} & a_{34} \\ a_{41} & a_{42} & a_{43} & a_{44} \end{bmatrix} \tag{47}$$

The error terms used follow the Cholesky decomposition. The upper triangle follows eq.(48). This order makes sense, as a negative productivity shock would also show up as a drop in weekly hours worked, then as an increase in unemployment as hiring slows and layoffs increase, and finally as a reduction in borrowing and lending. Weekly hours worked increased unemployment as firms under pressure reduced headcount and borrowing, while households became more nervous about the economy. Unemployment rates affect borrowing and lending as agents are less willing and able to take out loans.

$$
\begin{bmatrix} e_{1t} \\ e_{2t} \\ e_{3t} \\ e_{4t} \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ b_{21} & 1 & 0 & 0 \\ b_{31} & b_{32} & 1 & 0 \\ b_{41} & b_{42} & b_{43} & 1 \end{bmatrix} \begin{bmatrix} u_{1t} \\ u_{2t} \\ u_{3t} \\ u_{4t} \end{bmatrix} \tag{48}
$$

Recursive residual bootstrapping is used to create the confidence intervals. Recursive residual bootstrapping begins with fitting the VAR model and finding the residuals. From there, the residuals are resampled randomly to create synthetic data. The model is then refit with the synthetic data. This process is repeated a handful of times before an IRF is generated using the standard model, along with the 5th and 95th percentile values of the synthetic IRF for each time period.

Impulse response functions by running a sequence of OLS regressions for different horizons (h). The outcome in the future is based on a constant α_h plus the coefficient β_h of the shock to labor productivity, plus the covariates of the control variables $\gamma_h x_t$, plus the error terms v_{t+h} as seen in eq. (49). This unstructured method allows for flexibility without formally structuring the model. It is important to point out that the error term is structured v_{t+h} as a time series moving average of each of the errors that occurred before it as seen in eq. (50). This structure is based on the idea that a time series can be expressed in its infinite MA form, which helps maintain independence across horizons. Conditional Mean independence is preserved because the error term is assumed to be independent of the structural shock, conditional on control variables. The conditional mean independence ensures that the OLS regression of yields a consistent estimate. So, the MA structure for the errors helps maintain this independence across horizons, because MA processes have no persistent correlation with past shocks beyond their own lag structure.

$$
y_{t+h} = \alpha_h + \beta_h s_t + \gamma_h x_t + v_{t+h} \tag{49}
$$

$$
v_{t+h} = u_{t+h} + \psi_1 u_{t+h1} + ... + \psi_h ut \tag{50}
$$

7.4 Empirical Results

The VAR is computed for the combined data from all countries. The countries include: Austria, Belgium, Czechia, Denmark, Finland, France, Germany, Greece, Hungary, Ireland, Italy, Poland, Portugal, Spain, Sweden, and Norway. It is important to note that several countries were excluded from the original empirical model because they lacked household debt data, and Switzerland was specifically excluded due to very limited cyclical data from 2015 to 2019. The IRF in Figure (28) has a negative productivity shock of roughly -1.6% representing the standard deviation. It is visibly clear that the shock follows an AR(1) based on the restrictions that were imposed earlier.

The decrease in hours is not significant and is limited to roughly a 0.015 drop in weekly hours worked in the 3rd quarter after the shock. The cycle unemployment rate increases significantly by 0.25% during the 3rd through 5th periods after the shock, then slowly declines. The real debt shape mirrors the move in the unemployment rate, but instead shows a decrease in real debt. This shows that real debt drops by roughly 0.005% relative to trend debt after a shock. This margin is significant and rather persistent. This VAR has the 5 year Hamilton filter. The results of the 2-year filter show a similar framework, but are not as clean. This can be seen in the appendix as Figure (40)

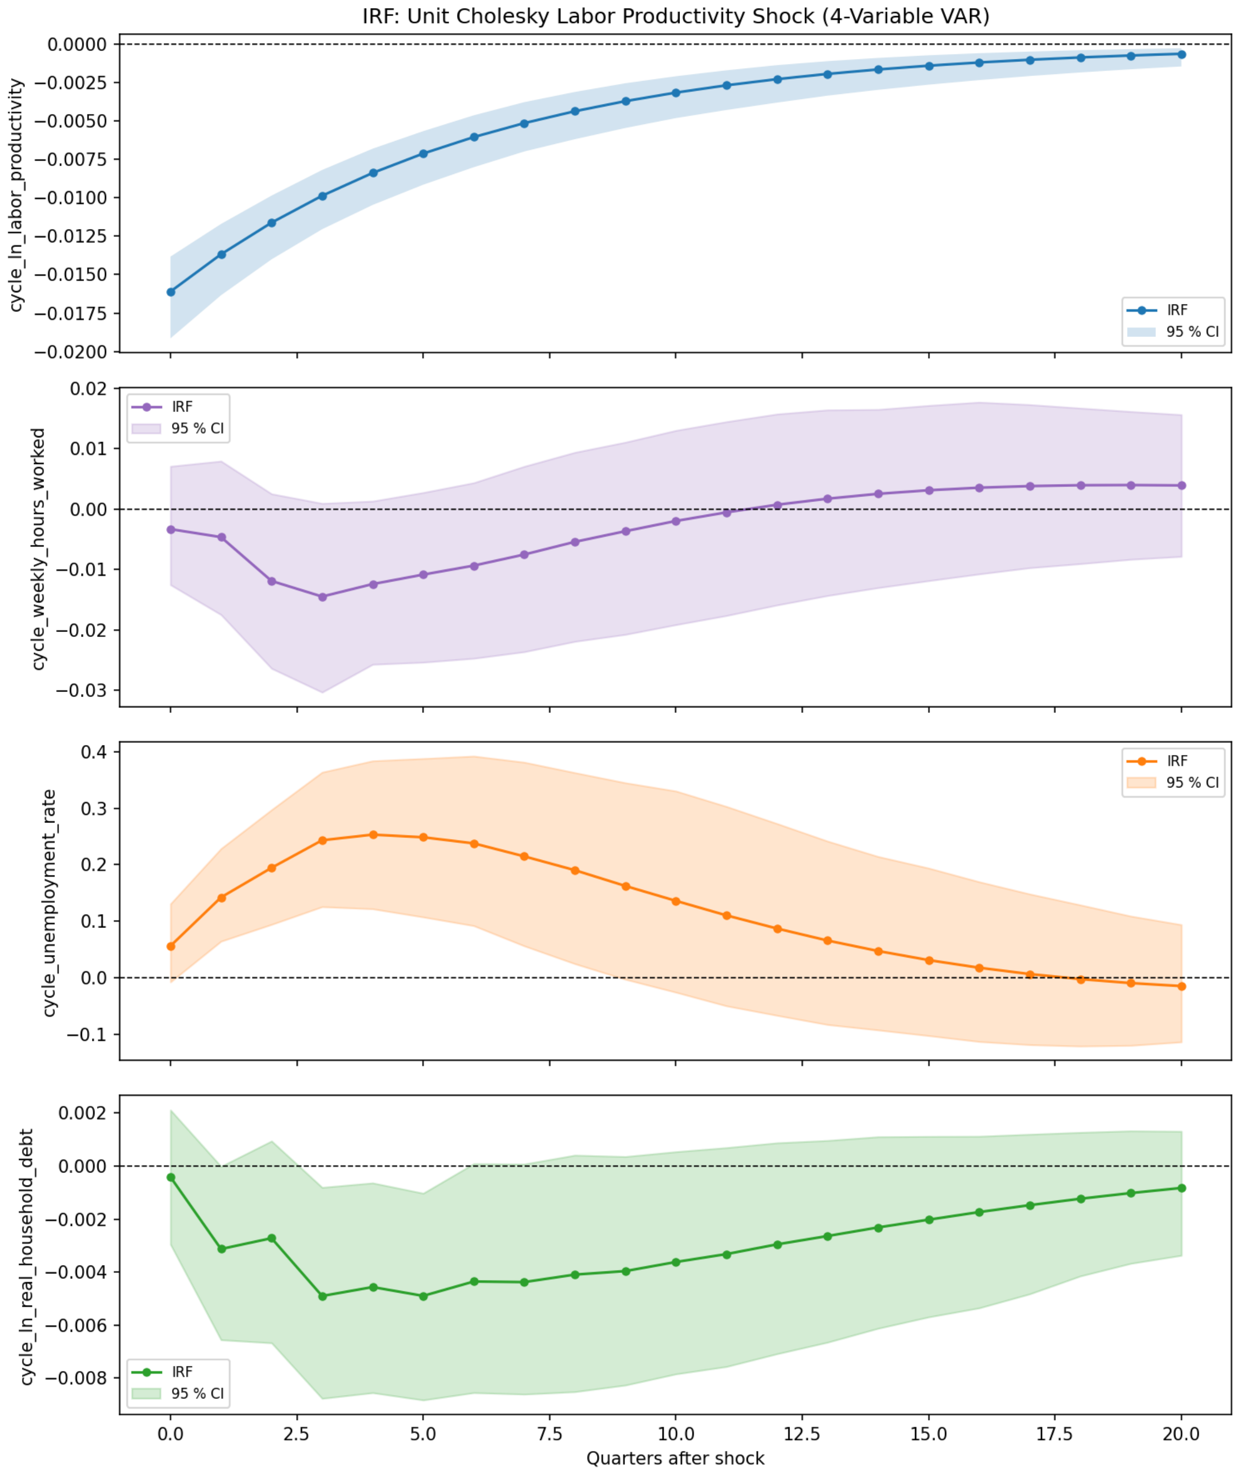

Figure 28: VAR All Countries

Figure (29) represents the LP's IRF to the productivity shock. As you can see, overall variable trends show similar directions and significance in the variables' movements, but with more chop. The labor productivity shock response is now flexible and shows a linear increase back to the starting level. Weekly hours worked show a small, rather insignificant decrease, except for the 4th quarter after the shock. The unemployment rate increases by roughly .25%, a similar amount as the VAR, before dropping back to the starting point. Real household debt decreases linearly until peaking at roughly 0.007% in the 10th quarter after the shock, and does so significantly. It is important to note that the two year Hamilton filter was executed and posted in the appendix as Figure (41). It shows the same general movements across all variables, but now real household debt decreases significantly only in the 8th period.

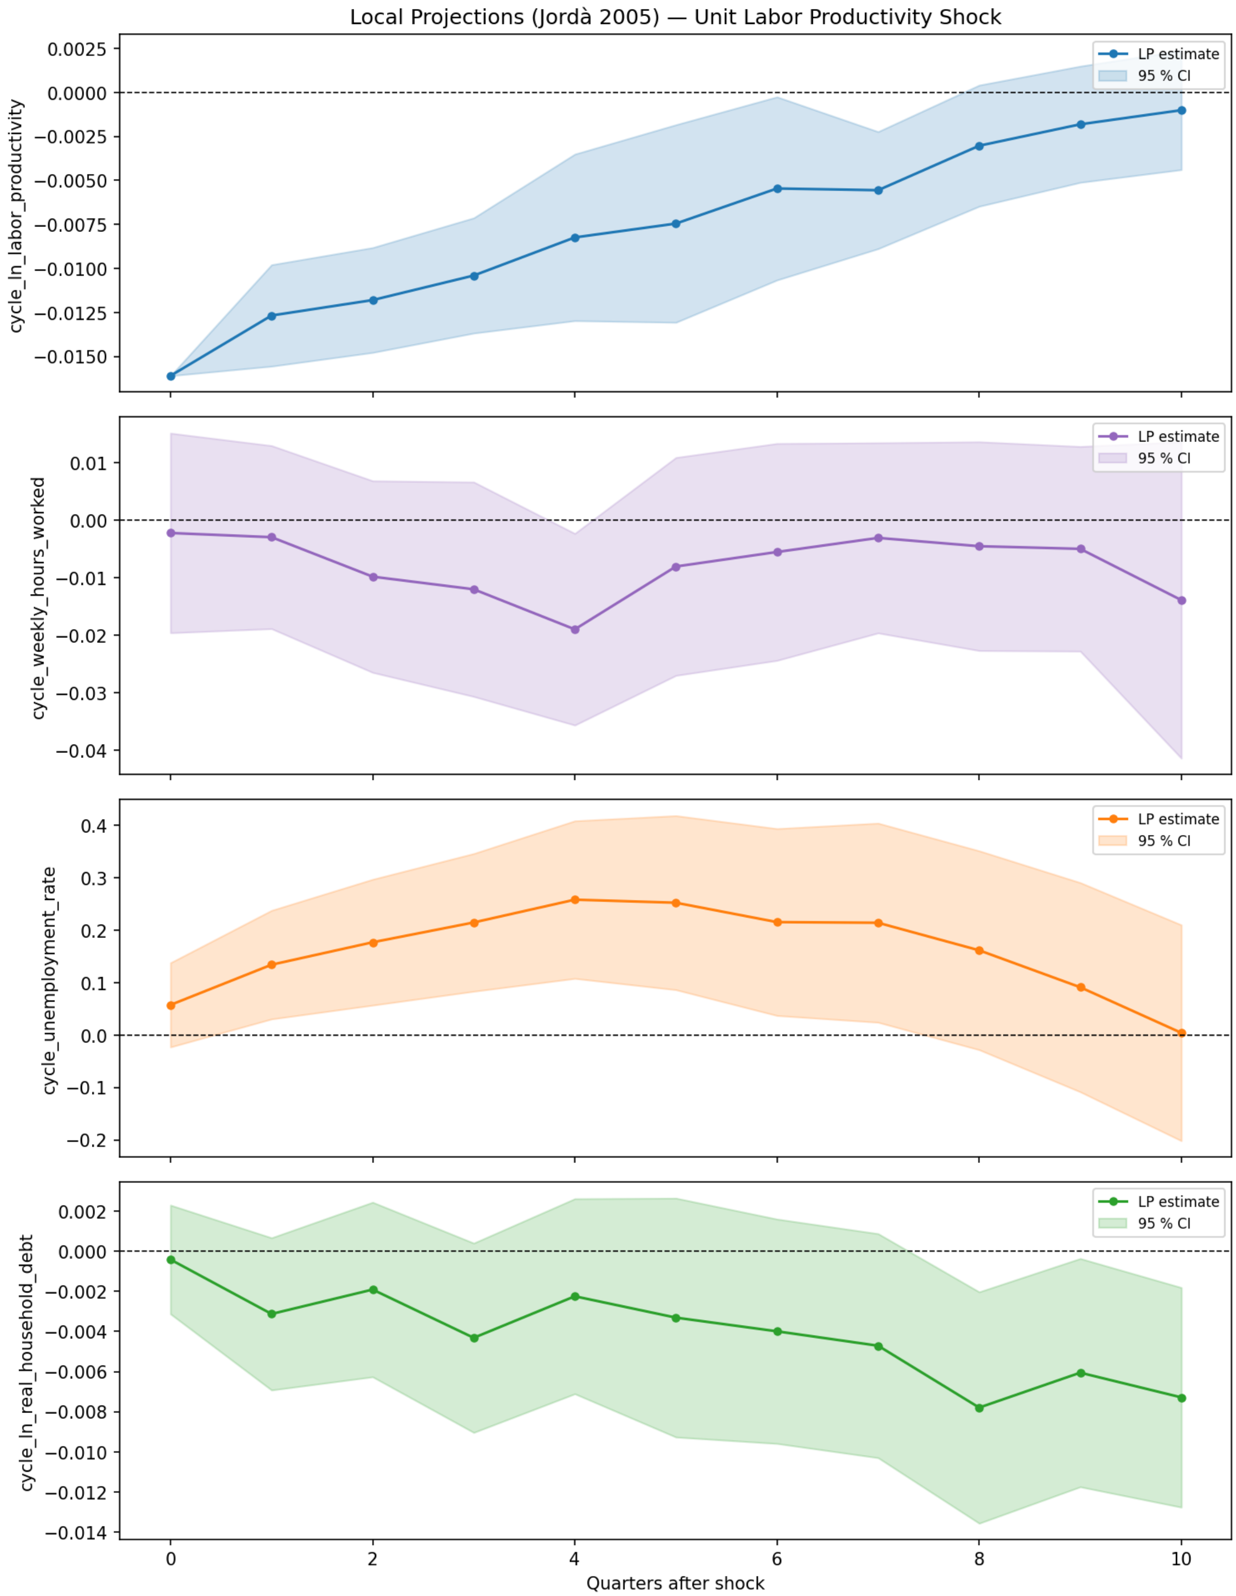

Figure 29: LP All Countries

From here on out, the VAR will be run for different splits based on high vs. low levels of variation in weekly hours worked, unemployment, real household debt, and the household debt-to-GDP ratio. This will show how countries with different dynamics have their debt perform differently. This can be seen in Figure (30). The countries are listed alphabetically and the standard deviation and average household debt to GDP values are displayed. The values are bolded if the countries value is greater than the median value.

Country Summary: Debt/GDP, Debt, Hours & Unemployment Volatility
(Bold = high group for that column; country names never bold)

Country	Hours Cycle Std Dev	Unemployment Std Dev	Debt Cycle Std Dev	Avg Debt/GDP
Austria	**0.8664**	0.5812	0.0381	50.1036
Belgium	0.2472	0.8772	0.0524	49.8333
Czechia	**0.3640**	**1.9170**	**0.1540**	21.4750
Denmark	**0.3955**	1.4211	0.0464	**107.8405**
Finland	0.1295	0.8903	**0.0716**	50.5102
France	**0.3063**	0.5749	0.0300	**51.5750**
Germany	0.1755	0.2436	0.0193	**56.9417**
Greece	0.2386	**6.4267**	**0.0993**	45.8839
Hungary	0.2096	**1.6171**	**0.1390**	22.6417
Ireland	**0.4253**	**2.6542**	0.0437	**78.5722**
Italy	**0.4108**	**1.4799**	**0.0648**	34.8807
Norway	0.1519	0.6012	0.0188	**78.4737**
Poland	0.1902	**2.3402**	**0.1121**	26.6934
Portugal	**0.2484**	**3.0464**	**0.0974**	**74.2989**
Spain	**0.2476**	**4.9537**	**0.1034**	**65.5795**
Sweden	0.1211	0.8527	0.0258	**71.6132**

Figure 30: Country Split Overview

Weekly hours worked were split between countries with high and low variation. The countries with high volatility in weekly hours are listed with their standard deviation: Austria (0.8664), Czechia (0.3640), Denmark (0.3955), France (0.3063), Ireland (0.4253), Italy (0.4108), Portugal (0.2484), and Spain (0.2476). The countries with low weekly variation are: Belgium (0.2472), Finland (0.1295), Germany (0.1755), Greece (0.2386), Hungary (0.2096), Norway (0.1519), Poland (0.1902), and Sweden (0.1211).

Figure (31) displays the VAR split. Both high and low volatility countries experience a similarly sized labor productivity shock. The low weekly hours worked volatility in countries has a more persistent shock. Looking at weekly hours, high volatility shows a large decrease of roughly 0.03 hours, while low volatility in weekly hours in countries has almost no effect. This is driven by the split itself. Going to the unemployment rate, we start to see differences. Both splits exhibit a similar path from the shock to 2 quarters after the shock. From there, the low hours volatility countries continue to rise before slowly decreasing. This is likely due to the larger persistence of the labor productivity shock. In countries with high hours volatility, the unemployment rate drops and even overshoots to lower unemployment before approaching the steady state. This leads to an interesting IRF for real household debt. Both splits drop to roughly the same.006% decrease in real debt, but the variance split does it sooner at the 4th quarter after the shock vs the 5th quarter of the low volatility shock. From here, the highly volatile country's debt returns to its original level fairly rapidly, while the low hours volatility countries take longer to recover. This is likely due to the unemployment rate staying higher for longer. This is a great finding, as the group of countries with high weekly hours volatility includes Portugal and Spain, just above the cutoff. They have a very high unemployment rate, high volatility, and volatile household debt. This means that the high weekly hours volatility includes many countries with high overall unemployment volatility, while the low hours volatility countries include those with a range of unemployment volatility and low hours volatility. This does a good job of highlighting that weekly hours volatility across countries does not meaningfully affect on real household leverage.

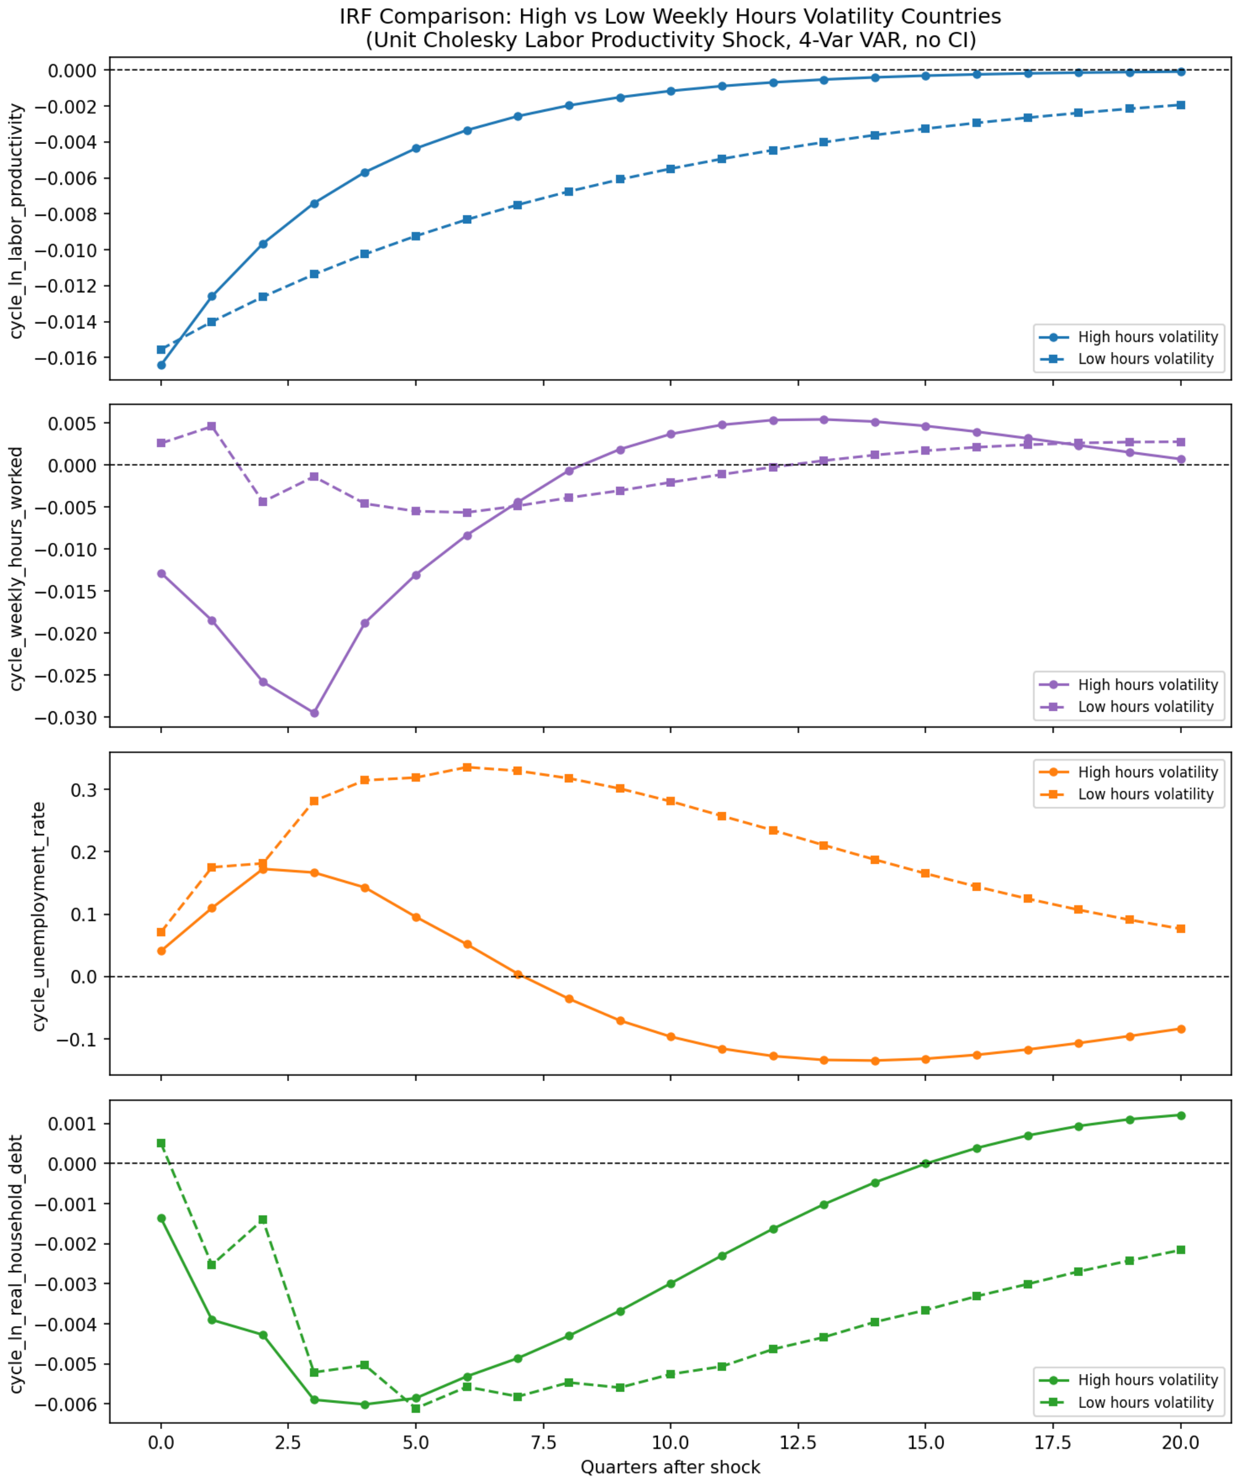

Figure 31: VAR High vs Low Weekly Hours Volatility Split

Next, we move on to the high vs low employment volatility split. The splits can be seen in Figure (32) The countries with high unemployment volatility are listed along with their standard deviations: Czechia (1.9170), Greece (6.4267), Hungary (1.6171), Ireland (2.6541), Italy (1.4799), Poland (2.3402), Portugal (3.0464), and Spain (4.9537). The low unemployment volatility countries include: Austria (0.5812), Belgium (0.8773), Denmark (1.4211), Finland (0.8903), France (0.5748), Germany (0.2435), Norway (0.6012), and Sweden (0.8527).

The shock to labor productivity is roughly twice the size for the high unemployment volatility countries, and the persistence of the high unemployment countries is less. Hours worked drop for the high unemployment volatility split and a return to normal, a standard feature. The low-volatility feature leaves weekly hours worked roughly unaffected until around 5 quarters after the shock, after which they begin to increase. This is a bit odd, as hours worked should not increase like this. Unemployment response has a similar shape for both the high and low splits. The main difference is that the magnitude of the response is roughly twice the size for the high unemployment volatility split. Real household debt decreases by a larger margin under the high-unemployment volatility benchmark and peaks at roughly 0.6% 3 periods after the shock. The low unemployment volatility split bottoms out 9 quarters after the shock, with a decrease of roughly 0.55%. This is extremely interesting, as it is happening with higher hours worked and a similar unemployment rate to that seen in the high split.

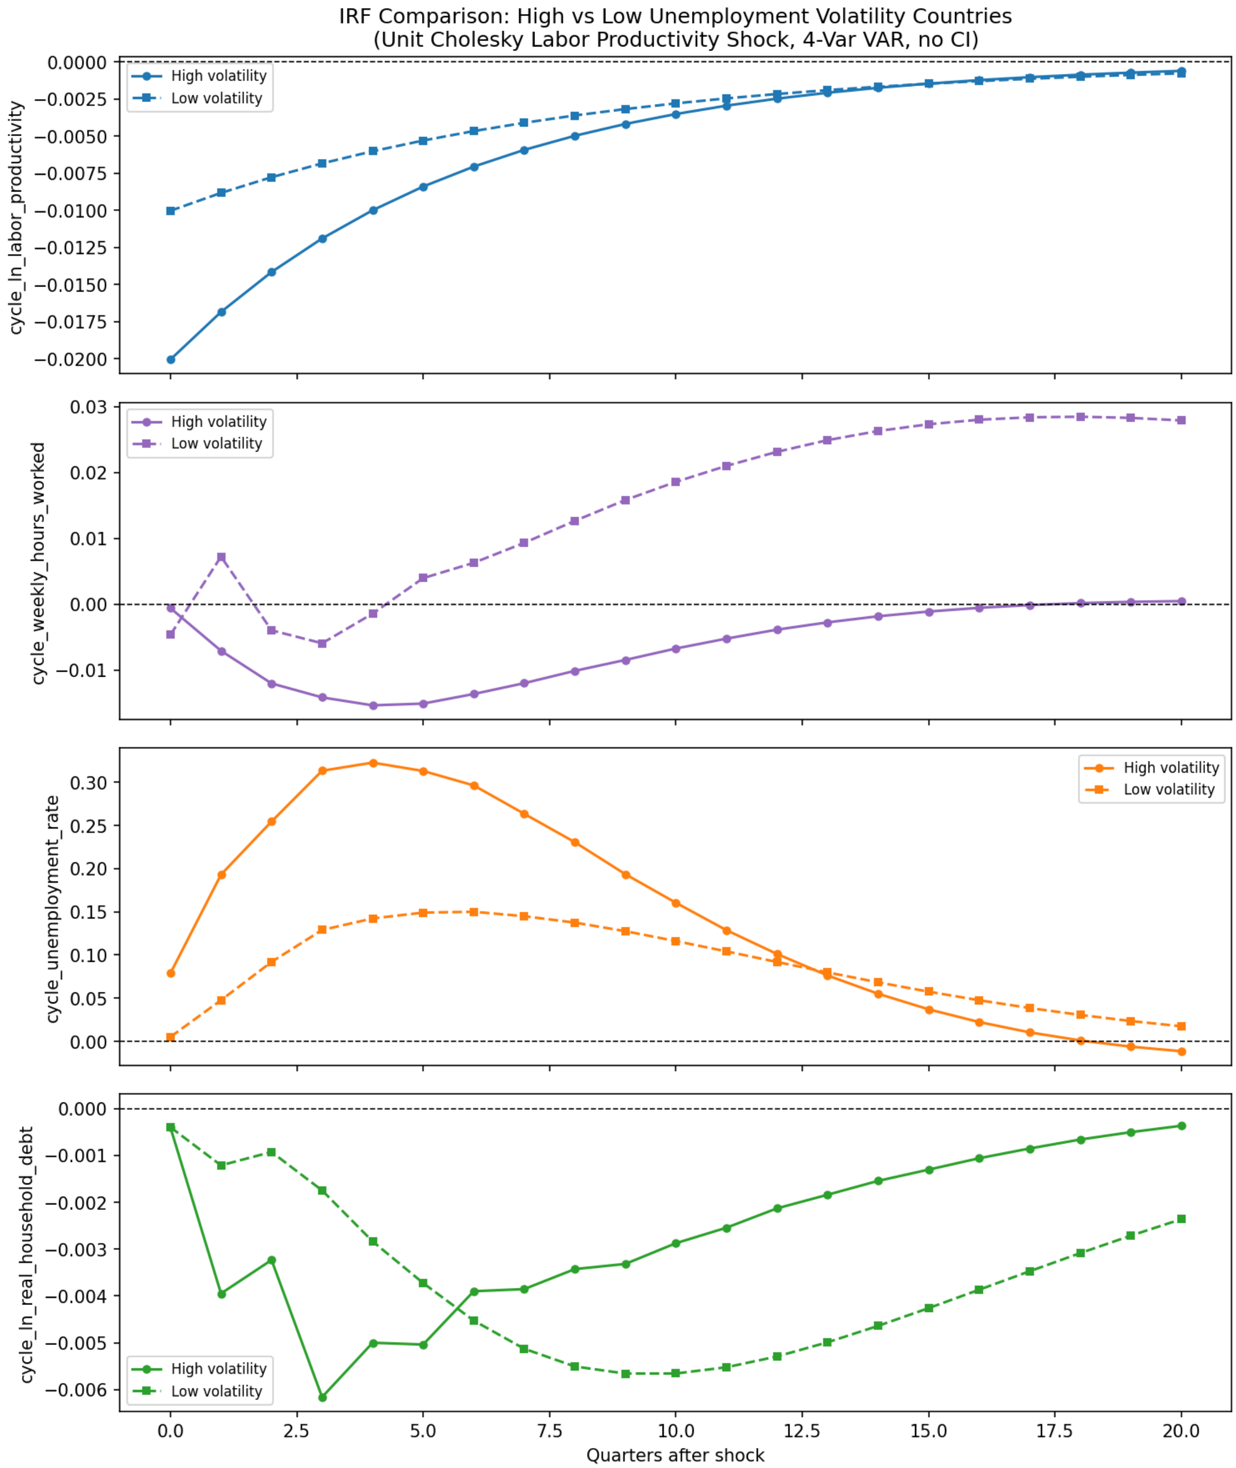

Figure 32: VAR High vs Low Unemployment Volatility Split

The next split of focus is the high vs low debt volatility splits. The countries with high volatility in real debt are listed with their standard deviation: Czechia (0.1540), Finland (0.0716), Greece (0.0993), Hungary (0.1390), Italy (0.0648), Poland (0.1121), Portugal (0.0974), and Spain

(0.1034). The countries with low volatility in real debt are listed with their standard deviation: Austria (0.0381), Belgium (0.0524), Denmark (0.0464), France (0.0299), Germany (0.0193), Ireland (0.0437), Norway (0.0188), and Sweden (0.0258).

Figure (33) presents the splits finding. This split specifically tends to tell a better story of what is happening and the underlying mechanics. First off, the labor productivity shock sizes are similar in magnitude, but the low debt volatility shock is slightly larger. The difference is that countries with high debt volatility have much longer persistence in their shocks. Weekly hours worked have a small response from both splits. The low-debt split initially drops but then increases to a positive response after the 5th period, following the shock. The high debt split starts positively in the first two periods, then drops slightly below the equilibrium. The unemployment response starts to paint the picture of why labor productivity persistence matters. The high debt volatility split has its unemployment rate increase by roughly 3x the magnitude of the low debt split, and one period later. The unemployment rate returns to its starting value 12 periods after the shock for the low-debt-volatility split and 20 quarters after the shock for the high debt volatility split. This mirrors what happens to the labor response when persistence of the shock increases in the HA model. The response to real debt differs significantly due to the split in real debt volatility. The response is fairly similar for both splits, starting at the shock point until 2 quarters after the shock. From here, the low debt volatility group has debt volatility increase to the starting value 5 periods after the shock and roughly stays there. The high debt split decreases further, to roughly a 0.7% decrease in leverage. Debt does not return to the steady state in the time frame and, in fact, is roughly half the maximum shock 20 quarters after the shock.

Recapping this finding, the shock sizes are fairly similar, but the high-debt variation split has much higher persistence. This doesn't show much change in hours worked, but we start to see large changes in unemployment and real household debt. This fits the story that firms face some kind of adjustment cost and countries that have higher persistence to their shocks tend to have their unemployment rate adjust by a larger magnitude and households within those countries make adjustments to their borrowing and saving because not only is the shock bigger, but they themselves see a higher and more persistent unemployment rate.

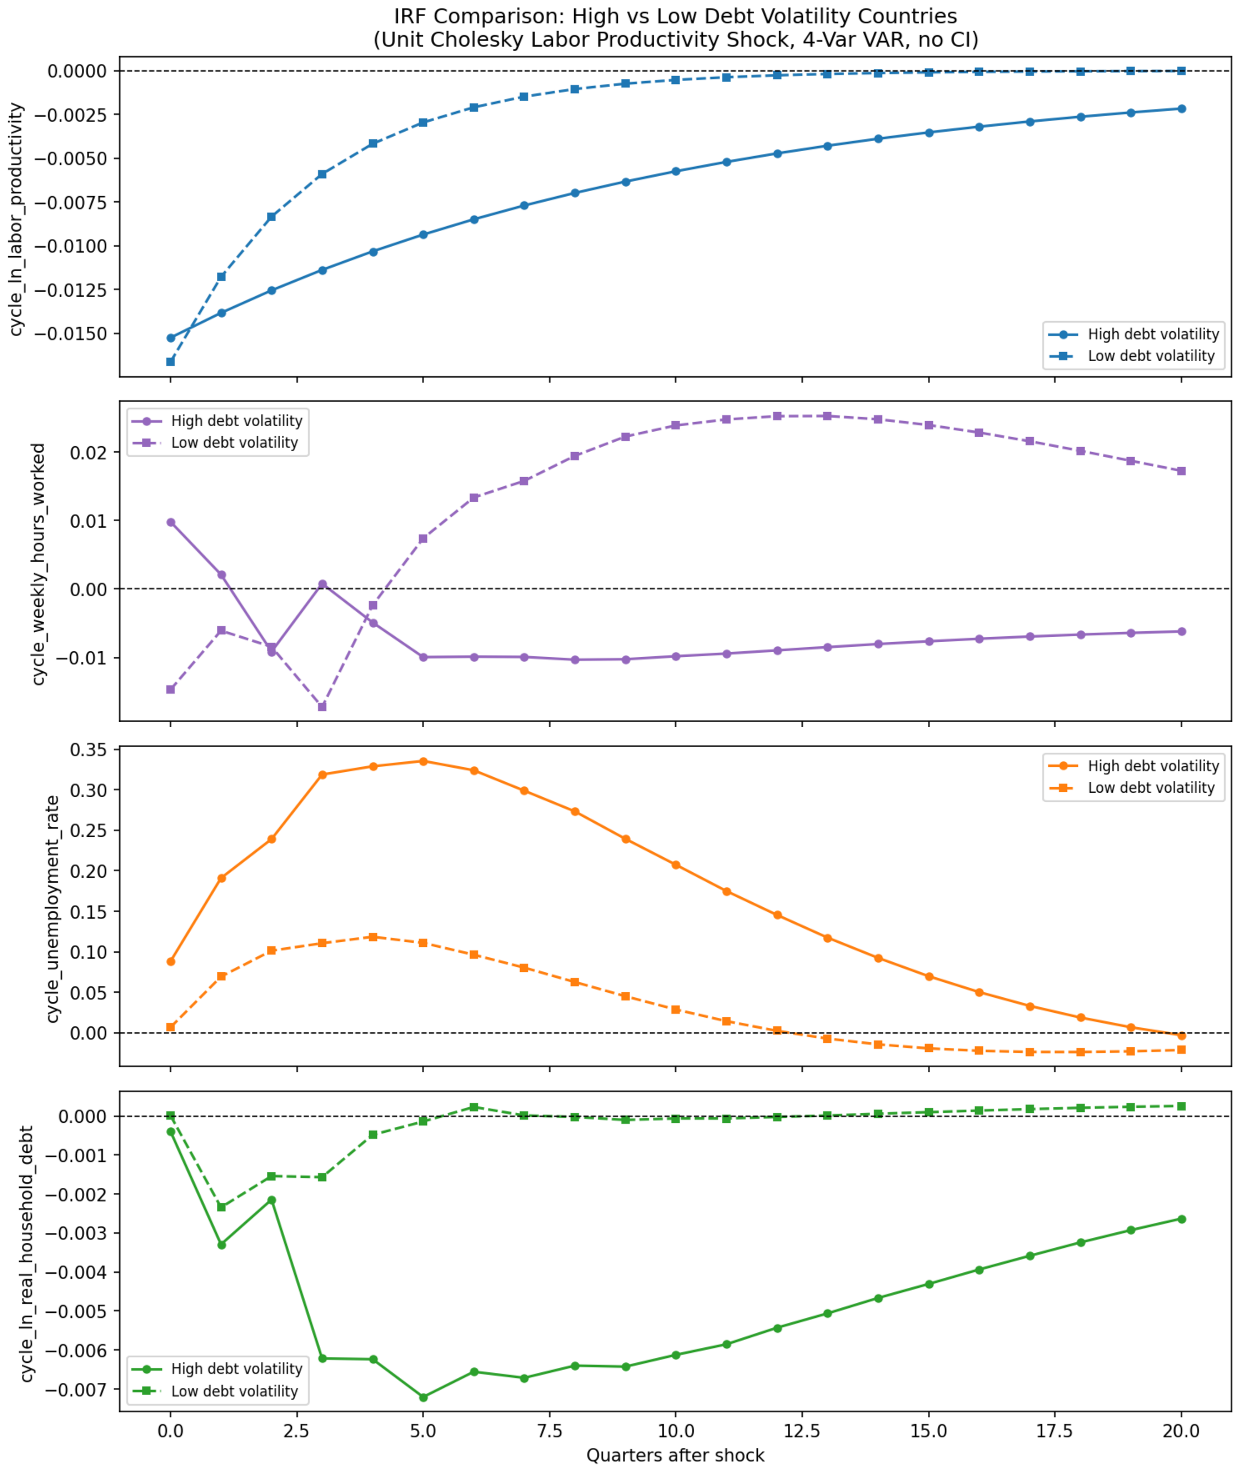

Figure 33: VAR High vs Low Debt Volatility Split

The countries were then split based on having high or low average household debt to GDP.

High average household debt to GDP countries are listed with mean debt/gdp: Denmark (107.8), France (51.6), Germany (56.9), Ireland (78.6), Norway (78.5), Portugal (74.3), Spain

(65.6), and Sweden (71.6). The low debt to GDP countries include: Austria (50.1), Belgium (49.8), Czechia (21.5), Finland (50.5), Greece (45.9), Hungary (22.6), Italy (34.9), and Poland (26.7).

Figure (34). Impulse Response Analysis by average debt levels. The findings in this section support much of what was found for the split on high vs low debt volatility. First, labor productivity shocks are of comparable magnitudes across the two splits; however, high debt to gdp countries exhibit marginally larger initial impacts. Notably, high debt environments are characterized by greater persistence of productivity shocks, a critical observation made during the last split. In terms of weekly hours worked, responses remain relatively muted in both regimes. The high debt level countries' weekly hours initially contract, then recover to positive territory 4 periods after the initial shock. Conversely, the low debt split shows a rather persistent decrease in hours.

The unemployment response elucidates why labor productivity persistence warrants special attention. Low debt to GDP countries experience approximately double the increase unemployment relative to High debt counterparts, with peak effects occurring one year sooner and dropping back to the original unemployment level 5 periods after the shock. The low debt countries employment level never recovers 20 periods after the shock.

Real debt responses further illuminate regime-dependent outcomes following the debt level split. Both groups exhibit similar behavior during impact and two subsequent quarters before diverging. The lhigh debt group returns to baseline, maintaining stability thereafter. In contrast, the high-debt split experiences an additional leverage reduction to approximately -1%, with debt levels failing to revert within the observation window, persisting at roughly one-half of the maximum shock magnitude twenty quarters after impact.

These patterns replicate key features observed in the high debt volatility split. An interpretation of the debt level split is that countries with lower household debt to GDP have more persistent shocks. The more persistent the socks, the more volatility in labor productivity. Unemployment's movements are more volatile and persistent. This causes changes in household borrowing and saving to decrease leverage by a larger amount with persistence. This unemployment volatility that households feel is causing them to initially choose less leverage. Another way of saying this is that households in countries with greater volatility in labor productivity and unemployment not only experience larger changes in leverage but also choose to take on less debt. While this cannot be directly observed in HA model because the level of unemployment is fixed, by increasing the level of unemployment at the same output level and looking at how debt levels decrease to increased risk backs up these findings.

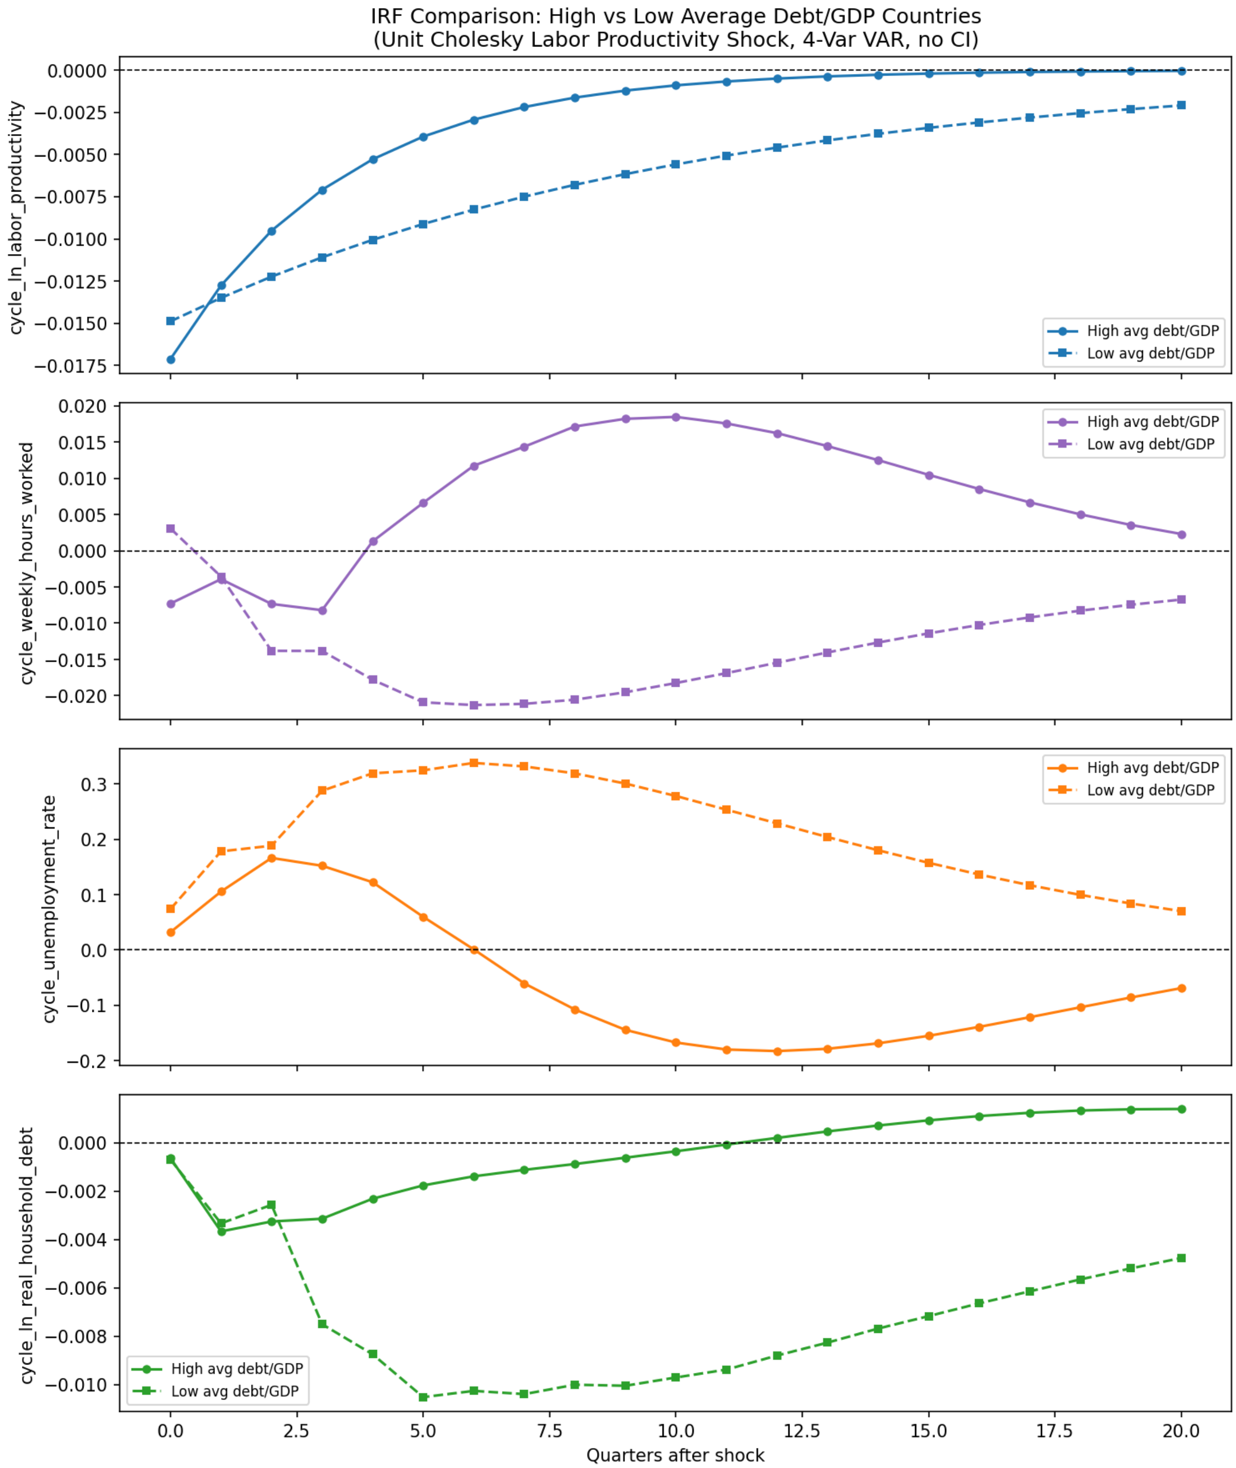

Figure 34: Var High vs Low Debt Level Split

7.5 Conclusion

Overall, the HA IRF does a decent job of capturing some of the dynamics. The aggregate VAR and LP show similar trends with a decrease in weekly hours worked, an increase in unemployment, and a decrease in real household debt. The key splits were high vs low real debt volatility and high vs low average household debt, which indicate that countries that split with a more persistent labor productivity shock experience larger and more prolonged increases in unemployment and a larger decrease in real household debt. These findings are significant and align with what occurs when the persistence of the labor productivity shock is increased in the HA model.

8 Conclusion

Most structural models rely on intensive margin adjustments (hours worked per agent) despite overwhelming empirical evidence for extensive-margin mechanisms. This paper resolves that gap by analyzing a three-period OLG model with exogenous youth unemployment and an endogenous Heterogeneous Agent Life cycle framework in which labor productivity drives transitions in employment status. Our comparative statics show that variations in the extensive margin, not just hours per worker, fundamentally reshape credit markets by lowering household debt levels. Impulse response function analysis of the Heterogeneous Agent (HA) model reveals that incorporating extensive margin adjustments yields economically meaningful deviations from conventional predictions. Specifically, variations in employment status yield distinct responses that lower both real interest rates and aggregate household debt. These dynamics are frequently overlooked by intensive margin focused models. Vector autoregressive and Local projection empirical tests validate these theoretical implications across countries, confirming that persistent labor productivity shocks generate larger, more prolonged increases in unemployment, leading to sharper and more prolonged declines in real household leverage.

Within the three period Overlapping Generations (OLG) framework, introducing youth unemployment exposes the loan demand transmission channel previously neglected by intensive margin models. Youth unemployment faces stricter borrowing constraints, as agents are less willing to lend to the unemployed. When the unemployment rate is changed exogenously, comparative statics show that the constraints depress aggregate loan demand, leading to lower real interest rates and a lower quantity of loans.

The Heterogeneous Agent framework enables a more nuanced characterization of household behavioral dynamics through several key extensions: age-dependent unemployment rate distributions, employment transition probabilities, dividend income, and precautionary savings motivations, all grounded in consumption savings decisions under employment uncertainty. Comparative statics are exercises that exogenously reduce labor productivity. This yields two distinct findings. First, there is a decrease in loan demand similar to the 3 period OLG model. Second, loan supply increases among agents with around zero assets, who exhibit precautionary saving behavior to buffer against heightened unemployment risk. The HA model is further analyzed in terms of its impulse response to a negative productivity shock. The response makes it clear that changes in unemployment rates can cause persistent household deleveraging dynamics. These findings contrast with canonical intensive margin models, which fail to capture heterogeneity in borrowing constraints and the heightened precautionary-saving motive arising from incomplete markets.

The dynamic responses of weekly hours worked, unemployment rates, and real household debt to labor productivity shocks is empirically analyzed using VAR and Local Projection (LP) methods. European panel data are employed due to limited U.S. state-level observations for these variables; this cross-country variation enhances identification beyond the limitations of aggregate U.S. data. Countries are grouped to assess heterogeneity in responses across weekly hours worked, unemployment, real household debt, and the household debt-to-GDP ratio. These analyses provide a direct evaluation of HA model channels and their predictive validity by demonstrating that labor productivity shocks drive aggregate hours worked primarily by raising unemployment rates. Additionally, country-grouping analysis reveals that nations exhibiting more persistent productivity shocks experience larger increases in both the duration and persistence of unemployment, which subsequently induce greater household debt deleveraging

However, the model presented has certain limitations. In particular, it does not include capital. In reality, increased loan supply (used for self-insurance) would likely lead to higher investment and capital accumulation, raising output beyond what the current HA steady state predicts. Still, this is complicated by adjustment costs, as seen in two-asset HA models, which discourage the use of illiquid assets, such as capital, for consumption smoothing. As a result, households are more likely to save as a precaution in liquid assets and to reduce investment in illiquid assets when faced with income risk. Furthermore, the model does not account for shopping externalities, which could further exacerbate downturns. In a setting where consumers can accept a disutility of shopping effort to purchase goods at lower prices, higher unemployment might raise the wage-to-price ratio because the unemployed have lower income and more time (lower disutility) to search for lower prices. This would depress prices. It is important to note that this would increase unemployment and cause further deleveraging, as firms receiving lower prices for goods would reduce labor demand further. This shift could lower the capital share of income and reduce investment, thereby exacerbating output declines.

This paper establishes the empirical validity of modeling unemployment as an endogenous channel. The Heterogeneous Agent model with the extensive margin has substantive implications for the evolution of both real household debt and real interest rates.

References

Achdou, Yves, Jiequn Han, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll. 2022. "Income and wealth distribution in macroeconomics: A continuous-time approach." *The review of economic studies* 89 (1): 45–86.

Aiyagari, S Rao. 1994. "Uninsured idiosyncratic risk and aggregate saving." *The Quarterly Journal of Economics* 109 (3): 659–684.

Attanasio, O, P Levell, H Low, and V Sánchez-Marcos. 2015. *Aggregating Elasticities: Intensive and Extensive Margins of Female Labour Supply. University College London.* Technical report. mimeo.

Auclert, Adrien. 2019. "Monetary policy and the redistribution channel." *American Economic Review* 109 (6): 2333–2367.

Bewley, Truman F. 1986. "Interest-bearing Money and the Equilibrium Stock of Capital." Unpublished manuscript, Yale University, written 1977–1986.

Challe, Edouard, Julien Matheron, Xavier Ragot, and Juan F Rubio-Ramirez. 2017. "Precautionary saving and aggregate demand." *Quantitative Economics* 8 (2): 435–478.

Chetty, Raj, Adam Guren, Day Manoli, and Andrea Weber. 2011. "Are micro and macro labor supply elasticities consistent? A review of evidence on the intensive and extensive margins." *American Economic Review* 101 (3): 471–475.

———. 2013. "Does indivisible labor explain the difference between micro and macro elasticities? A meta-analysis of extensive margin elasticities." *NBER macroeconomics Annual* 27 (1): 1–56.

Clarida, Richard, Jordi Gali, and Mark Gertler. 1999. "The science of monetary policy: a new Keynesian perspective." *Journal of economic literature* 37 (4): 1661–1707.

Deaton, Angus. 1991. "Saving and liquidity constraints." *Econometrica* 59 (5): 1221–1248.

Eggertsson, Gauti B, Neil R Mehrotra, and Jacob A Robbins. 2019. "A model of secular stagnation: Theory and quantitative evaluation." *American Economic Journal: Macroeconomics* 11 (1): 1–48.

Gagliardone, Luca, and Mark Gertler. 2023. "Oil prices, monetary policy and inflation surges."

Galí, Jordi, Frank Smets, and Rafael Wouters. 2012. "Unemployment in an estimated New Keynesian model." *NBER macroeconomics annual* 26 (1): 329–360.

Gourinchas, Pierre-Olivier, and Jonathan A Parker. 2002. "Consumption over the life cycle." *Econometrica* 70 (1): 47–89.

Guerrieri, Veronica, and Guido Lorenzoni. 2017. "Credit crises, precautionary savings, and the liquidity trap." *The Quarterly Journal of Economics* 132 (3): 1427–1467.

Hamilton, James D. 2018. "Why you should never use the Hodrick-Prescott filter." *Review of Economics and Statistics* 100 (5): 831–843.

Hansen, Gary D. 1985. "Indivisible labor and the business cycle." *Journal of monetary Economics* 16 (3): 309–327.

Haughwout, Andrew, Donghoon Lee, Daniel Mangrum, Joelle Scally, Wilbert Van der Klaauw, and Crystal Wang. 2024. "Delinquency Is Increasingly in the Cards for Maxed-Out Borrowers." *Liberty Street Economics.*

Heckman, James. 1984. "Comments on the Ashenfelter and Kydland papers." In *Carnegie-Rochester Conference Series on Public Policy,* 21:209–224. Elsevier.

Heckman, James J, and Thomas E MaCurdy. 1980. "A life cycle model of female labour supply." *The Review of Economic Studies* 47 (1): 47–74.

Huggett, Mark. 1993. "The risk-free rate in heterogeneous-agent incomplete-insurance economies." *Journal of economic Dynamics and Control* 17 (5-6): 953–969.

Imrohoroğlu, Ayşe. 1989. "Cost of business cycles with indivisibilities and liquidity constraints." *Journal of Political economy* 97 (6): 1364–1383.

Janko, Zuzana. 2008. "Nominal wage contracts, labor adjustment costs and the business cycle." *Review of Economic Dynamics* 11 (2): 434–448.

Jordà, Òscar. 2005. "Estimation and inference of impulse responses by local projections." *American economic review* 95 (1): 161–182.

Jordà, Òscar, Moritz Schularick, and Alan M Taylor. 2017. "Macrofinancial history and the new business cycle facts." *NBER macroeconomics annual* 31 (1): 213–263.

Kaplan, Greg, Kurt Mitman, and Giovanni L Violante. 2020. "The housing boom and bust: Model meets evidence." *Journal of Political Economy* 128 (9): 3285–3345.

Kaplan, Greg, and Giovanni L Violante. 2014. "A model of the consumption response to fiscal stimulus payments." *Econometrica* 82 (4): 1199–1239.

Kaplan, Greg, and Giovanni L Violante. 2018. "Microeconomic heterogeneity and macroeconomic shocks." *Journal of Economic Perspectives* 32 (3): 167–194.

Katz, Lawrence F. 1986. "Efficiency wage theories: A partial evaluation." *NBER macroeconomics annual* 1:235–276.

Krueger, Dirk, Kurt Mitman, and Fabrizio Perri. 2016. *On the distribution of the welfare losses of large recessions.* Technical report. National Bureau of Economic Research.

Krusell, Per, and Anthony A Smith Jr. 1998. "Income and wealth heterogeneity in the macroeconomy." *Journal of political Economy* 106 (5): 867–896.

Kydland, Finn E, and Edward C Prescott. 1982. "Time to build and aggregate fluctuations." *Econometrica: Journal of the Econometric Society,* 1345–1370.

Lechthaler, Wolfgang, and Dennis J Snower. 2013. "Quadratic labor adjustment costs, business cycle dynamics, and optimal monetary policy." *Macroeconomic Dynamics* 17 (2): 464–475.

McGrath, Maggie. 2016. "63% Of Americans Don't Have Enough Savings To Cover A $500 Emergency,"." *Forbes Magazine.*

McKay, Alisdair, and Ricardo Reis. 2016. "The role of automatic stabilizers in the US business cycle." *Econometrica* 84 (1): 141–194.

Meyer, Bruce D. 2002. "Labor supply at the extensive and intensive margins: The EITC, welfare, and hours worked." *American economic review* 92 (2): 373–379.

Quercia, Roberto G, George W McCarthy, and Susan M Wachter. 2003. "The impacts of affordable lending efforts on homeownership rates." *Journal of Housing Economics* 12 (1): 29–59.

Ravn, Morten O, and Vincent Sterk. 2017. "Job uncertainty and deep recessions." *Journal of Monetary Economics* 90:125–141.

Rendahl, Pontus. 2016. "Fiscal policy in an unemployment crisis." *The Review of Economic Studies* 83 (3): 1189–1224.

Rotemberg, Julio J, and Michael Woodford. 1997. "An optimization-based econometric framework for the evaluation of monetary policy." *NBER macroeconomics annual* 12:297–346.

Shimer, Robert. 2012. "Reassessing the ins and outs of unemployment." *Review of Economic Dynamics* 15 (2): 127–148.

Sims, Christopher A. 1980. "Macroeconomics and reality." *Econometrica: journal of the Econometric Society,* 1–48.

9 Appendix

9.1 Adding Labor Force Participation to Employment Variance Decomp

To address this question, I use data from Eurostat and the OECD from 2002 to 2019. Eurostat provides weekly hours worked, total hours worked within in country, and household debt. The OECD has data on nominal GDP and the unemployment rate. These datasets are well suited to the question because they offer consistent, cross-country data across European nations. The data will be used to decompose the variance in total hours worked into the percent of variance caused by the employment rate, weekly hours worked, labor force participation, and the covariances between terms. The data will also be useful for analyzing labor market fluctuations and their relationship to household leverage. The focus on Europe is deliberate. The presence of a shared central bank and common monetary policy, along with high capital mobility across EU countries, helps isolate the role of loan demand caused by fluctuations in employment rather than interest rate fluctuations in driving household debt-to-GDP dynamics. This setup allows for a clearer identification of how employment outcomes relate to household borrowing behavior in a coordinated financial environment.

The variance decomposition of total hours worked follows the methodology outlined in (Hansen 1985). To extract cyclical components, all series are detrended using the Hodrick-Prescott (HP) filter with a smoothing parameter of 100, which is appropriate for the annual frequency data. The detrended value is then calculated as the ratio of the detrended (cyclical) component to the trend, plus one, a formulation that facilitates interpretation in log-linear terms. The variances and covariances are computed using the natural log of the detrended value of each term. It is important to note that covariances can be negative if the underlying variables are negatively correlated. The variance decomposition following equation (1). T stands for total hours worked in an economy over the past year. E stands for the employment rate. LF is the labor force size. H stands for weekly hours worked. The decomposition is performed in an equal-weighted manner across all data points to construct aggregate measures. Additionally, a within-country variance decomposition is conducted for each country in the sample.

$$\sigma^2_{\ln T} = \sigma^2_{\ln E} + \sigma^2_{\ln LF} + \sigma^2_{\ln H} + 2\,\sigma_{\ln E,\ln LF} + 2\,\sigma_{\ln E,\ln H} + 2\,\sigma_{\ln H,\ln LF}.$$

To examine the relationship between labor market dynamics and household indebtedness, detrended employment and detrended household debt-to-GDP ratios were analyzed across European countries. A two-year lag in detrended employment was selected based on empirical fit, reflecting the realistic delay with which households respond to labor market conditions by adjusting their borrowing behavior, both in leveraging during periods of high employment and deleveraging during periods of lower employment. The correlation is then conducted in an equal-weighted fashion across all available country-year observations to construct aggregate patterns. Within-country was assessed to capture the heterogeneity in the employment to household debt to GDP correlation within each country.

The aggregate variance decomposition was computed in an equal-weighted manner across all data points to ensure that no single country disproportionately influences the aggregate results. This provides a more balanced view of overall labor market dynamics within Europe, as different countries dealt with different idiosyncratic shocks. This approach allows for the construction of aggregate measures that reflect the central tendencies of the countries as a whole, rather than being skewed by larger economies. Additionally, within-country decompositions were done to show important heterogeneity across nations, enabling the identification of distinct labor market patterns that may be masked in the aggregate. Together, these two levels of analysis reinforce the robustness of the findings by showing that the significance of the extensive margin is not driven by outliers and holds consistently across multiple contexts.

To strengthen my analysis of the relationship between household debt-to-GDP and employment, I first examine detrended employment and detrended household debt-to-GDP at the country level. It indicates that there is a correlation, and the choice of using lagged employment deviation. Specifically, a two-year lag aligns empirically and intuitively with the time it takes for households to adjust their borrowing behavior in response to labor market changes. Once the country-level patterns are established, an aggregate correlation is constructed by pooling all countries in an equal-weighted manner, ensuring no single economy skews the results. This dual approach within the country and in the aggregate strengthens the robustness of the findings by capturing both aggregate and within-country dynamics.

In Figure (35), the aggregate variance decomposition of total hours worked reveals that fluctuations in the employment rate are the primary driver of variation across countries at over 60 %. The second largest contributor is the variance in the labor force size, followed by the covariance between the employment rate and average weekly hours worked. The variance in weekly hours ranks fourth. This is significant because not only is the employment rate, extensive margin, variation larger than the intensive margin of hours worked, but so is the often-ignored deviation in the labor force size. The covariance between the employment rate and labor force participation contributes a smaller, though still positive, portion. Lastly, the covariance between weekly hours and the labor force is negative, reflecting the negative correlation between these two components. This analysis underscores the pivotal role of the extensive margin, driven by fluctuations in employment rates, in explaining variations in total hours worked across European countries.

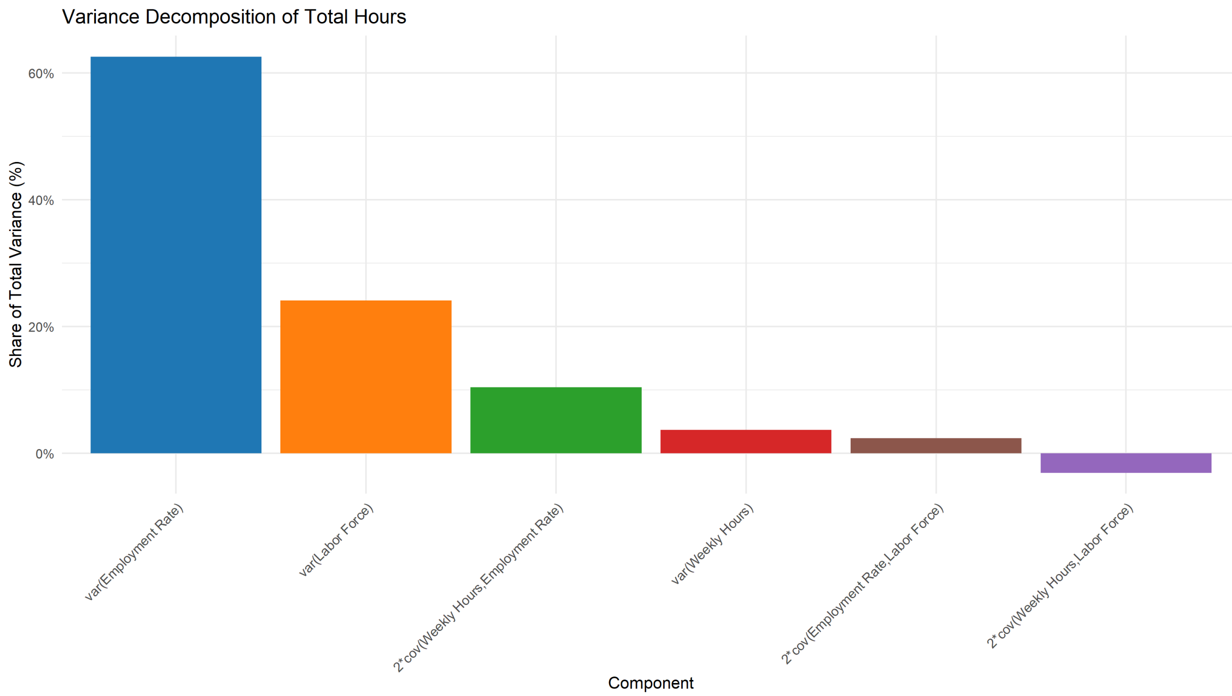

Figure 35: Aggregate Variance Decomposition

Austria, Denmark, France, Luxembourg, and Sweden stand out as exceptions to the broader trend observed across Europe, where the deviation in the employment rate is typically the dominant driver of total hours worked. Among these, Austria is unique in being primarily influenced by changes in weekly hours worked. This is due to Austria's relatively stable employment rate during the Great Recession, which only rose modestly from 4.4% in 2008 to 5.7% in 2009, and remained below 5.7% until 2014. This reflected a milder recessionary impact on the employment rate. Denmark, France, Luxembourg, and Sweden exhibit a variance in total hours worked that is predominantly driven by deviations in labor force size. These fluctuations stem from shifts in the native working-age population, migration patterns, and labor force participation rates. Despite these exceptions, the extensive margin of deviations in the employment rate was the primary contributor to labor market fluctuations in the majority of European countries, as illustrated in Figure (36).

Figure 36: Within Country Variance Decomposition

9.2 Heterogeneous Agent Alternative Specifications

Figure (37) shows the loan demand and loan supply if we add government debt (B gov) and government expenditures (G). Government Debt was set to 70% of GDP, and government expenditures were set to 18% of GDP; these are roughly in line with these values since the 1990s. The loan amount is roughly 1.4 because the debt is 0.7; that means household borrowing is roughly 0.7 household debt to GDP, and is much more in line with the actual data. As the interest rates rise, household borrowing is crowded out. The effects are roughly the same through both margins; they are just at higher levels of interest and total loan amount.

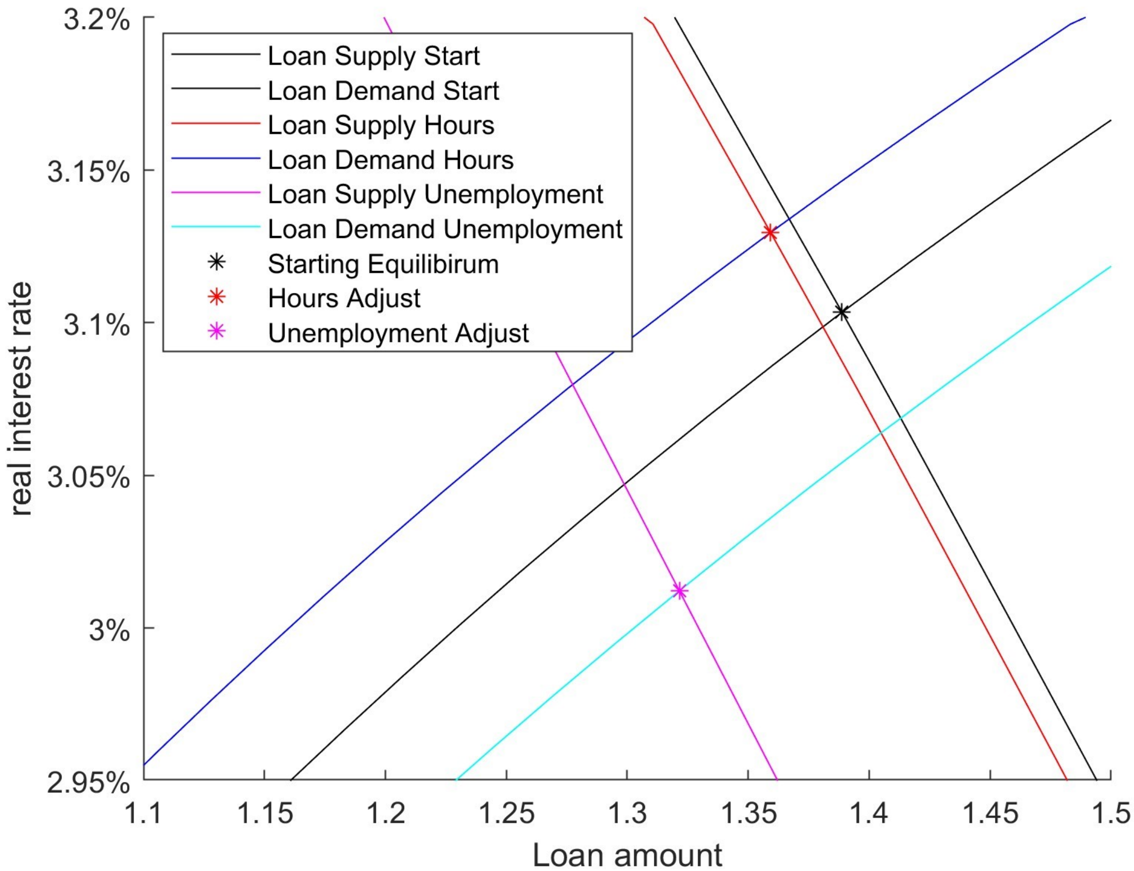

Figure 37: Het Loan Demand and Loan Supply with B gov and G

Figure (38) displays the intensive and extensive margin models with no smoothing. This produces very similar results to the smoothed version. The main difference is that the real interest rate jumps up and back down between 4 and 5 quarters after the shock. While this is slightly more erratic without smoothing, there are no significant differences between the unsmoothed and smoothed versions.

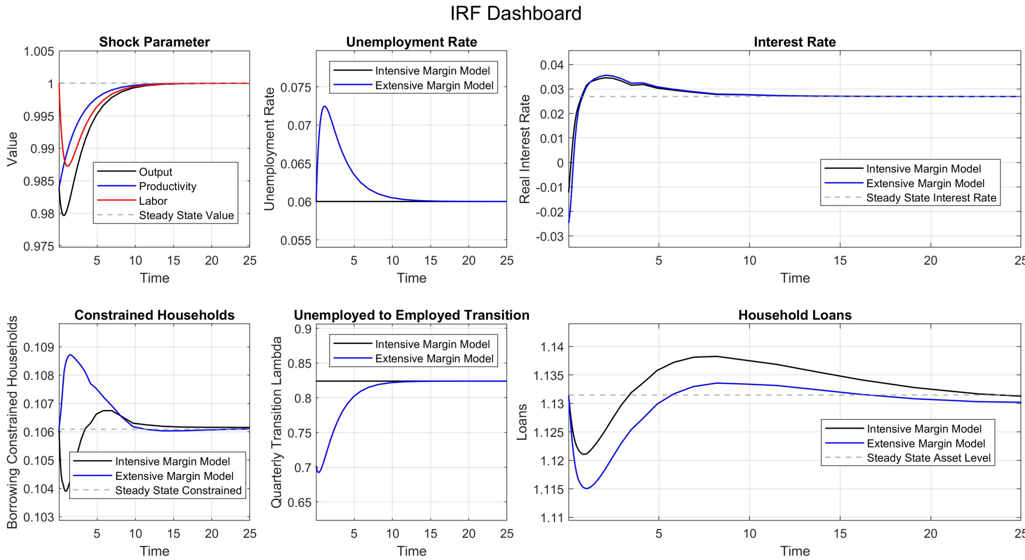

Figure 38: HA IRF with no smoothing

Figure (39) shows what would happen if you added government debt and deficits. One thing to note is that if the government has taken out debt, that means that the households are holding the debt and they have a higher asset level as a whole. A consequence of this is that the percent of constrained households is roughly 6.9% vs the 10.6 % when there is no government debt and government expenditures. Government debt was set to .7 or 70% of GDP. Government purchases of goods and services were set to 0.18, or 18%, of GDP. This is separate from interest payments and Unemployment/Social Security. Now the government market clearing condition is given by equation (51), which states that the revenue equals unemployment and social security transfers, plus interest payments, plus government expenditures, plus change in government debt. Government revenue will be increased by adjusting the tax rate. The change in government debt is governed by eq (52). The change in debt is interest expense, unemployment transfer payments, and government purchases of goods and services, less tax revenue that would be generated with the equilibrium interest rate, less the amount of tax revenue that would be collected if the tax rate was still the equilibrium rate, less a factor times the value of current government debt, less the equilibrium value. It is important to note that in order for the model to converge ϕ_g must be larger than the interest rate.

$$revIncome = expTransf + Bgov_t * r_t + G_t + \dot{Bgov}_t \tag{51}$$

$$B_dot = Bgov_t * r_t + expTransf + G_t - \sum_e \int_l \int_a \tau_{equm}(w_t(z_l h_{e,t}) + \pi_{l,t}) - \phi_g * (Bgov_t - Bgov_{equm}) \tag{52}$$

The IRF response shows a few things. First, the interest rate shock is not immediate and, in fact, takes longer to decrease before increasing and overshooting the equilibrium rate. It is important to note that when the rate is lower, government interest payments decrease, and when the rate is higher, they increase. Constrained households increase similarly under the extensive margin model, but then decrease and remain lower for a much longer time. Household debt decreases more slowly, but there is greater deleveraging because the process takes longer. From here, household debt takes longer to return to its starting level. It is important to note that a factor that contributes to this is that more lending flows to the government, which increases interest rates and causes a crowding out of household borrowing.

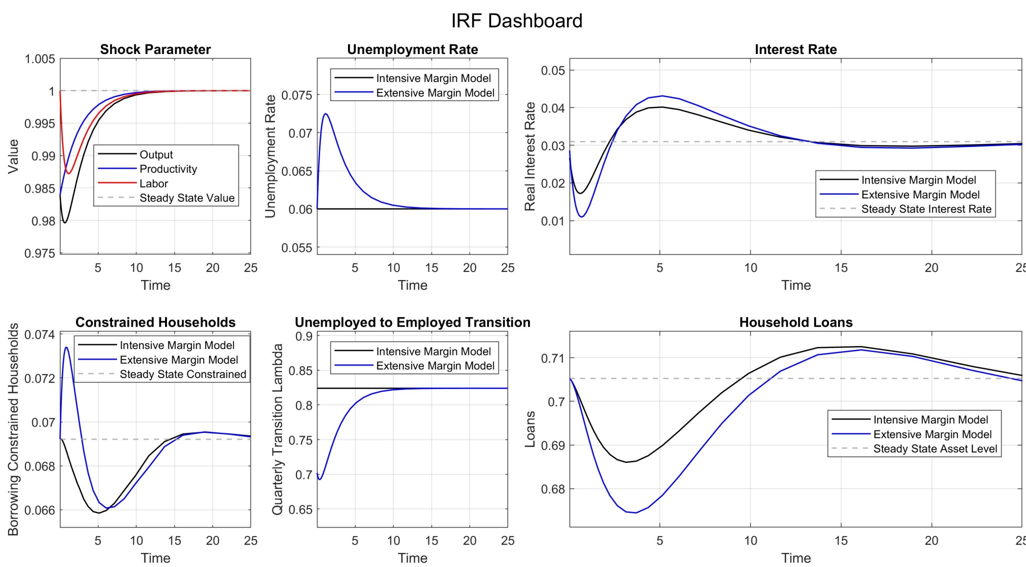

Figure 39: HA IRF with G Bgov and Deficits

9.3 Empirical IRF Additional Specifications

Figure (40) shows aggregate VAR for all of the countries with the Hamilton filter set to two. As you can see, the overall shapes are similar for the following: Labor productivity is a standard AR(1), weekly hours worked show no significant change, and unemployment increases by a statistically significant amount. The main difference here is that real household debt now increases over the first 3 periods before decreasing and returning to zero. There is now no significance to the real in any period. While this is not promising, this went somewhat as expected, as (Jordà, Schularick, and Taylor 2017) argued for looking ahead 5 years when examining debt cycles.

Figure (41) shows the LP with the Hamilton filter set to 2 years. The two year filter gives a rather similar result. The productivity shock follows the same linear approach back to normal: there is no significant movement in weekly hours worked, unemployment increases significantly before returning to normal, and real household debt starts at an insignificant value before decreasing linearly. It is important to note that the decrease in household debt is statistically significant only 8 quarters after the shock. Overall, both VAR and LP with the 2 year Hamilton filter show similar results for labor productivity shock, weekly hours worked, and the unemployment rate. The main difference is that the decrease in household debt is no longer statistically significant, as other papers have pointed out the need to use a longer filter period when examining debt leverage.

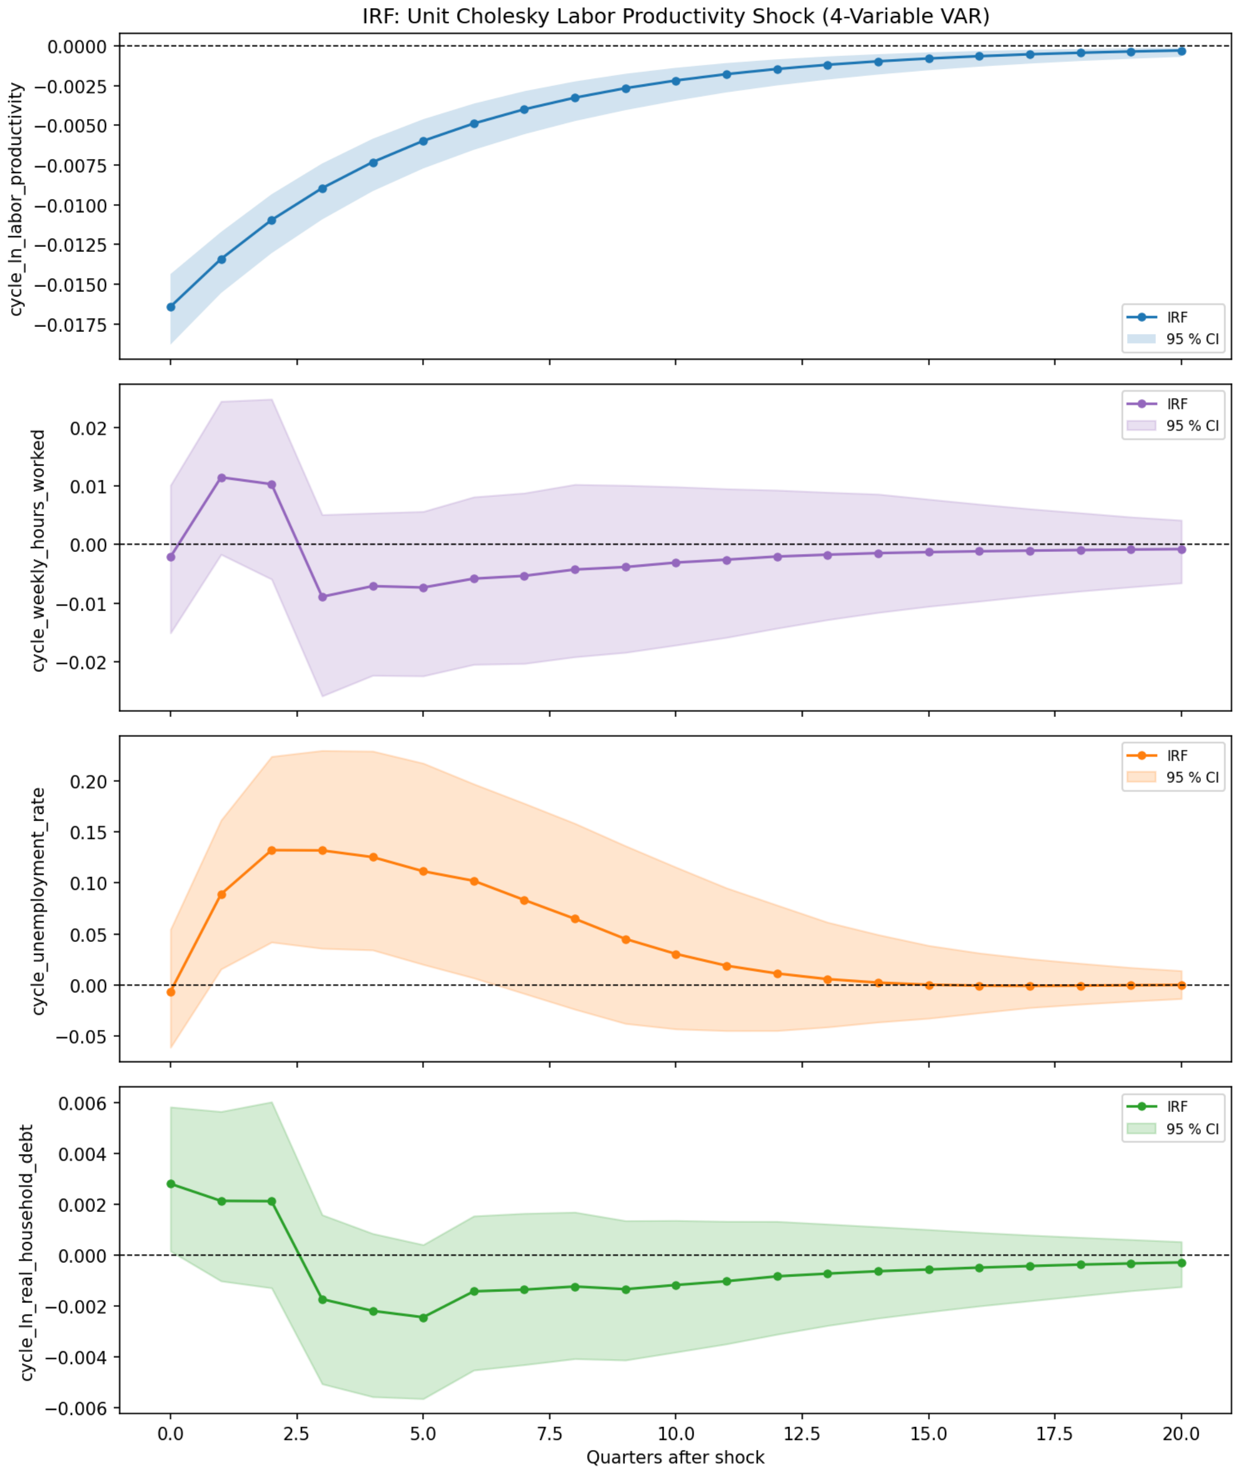

Figure 40: VAR All Countries 2 Year filter

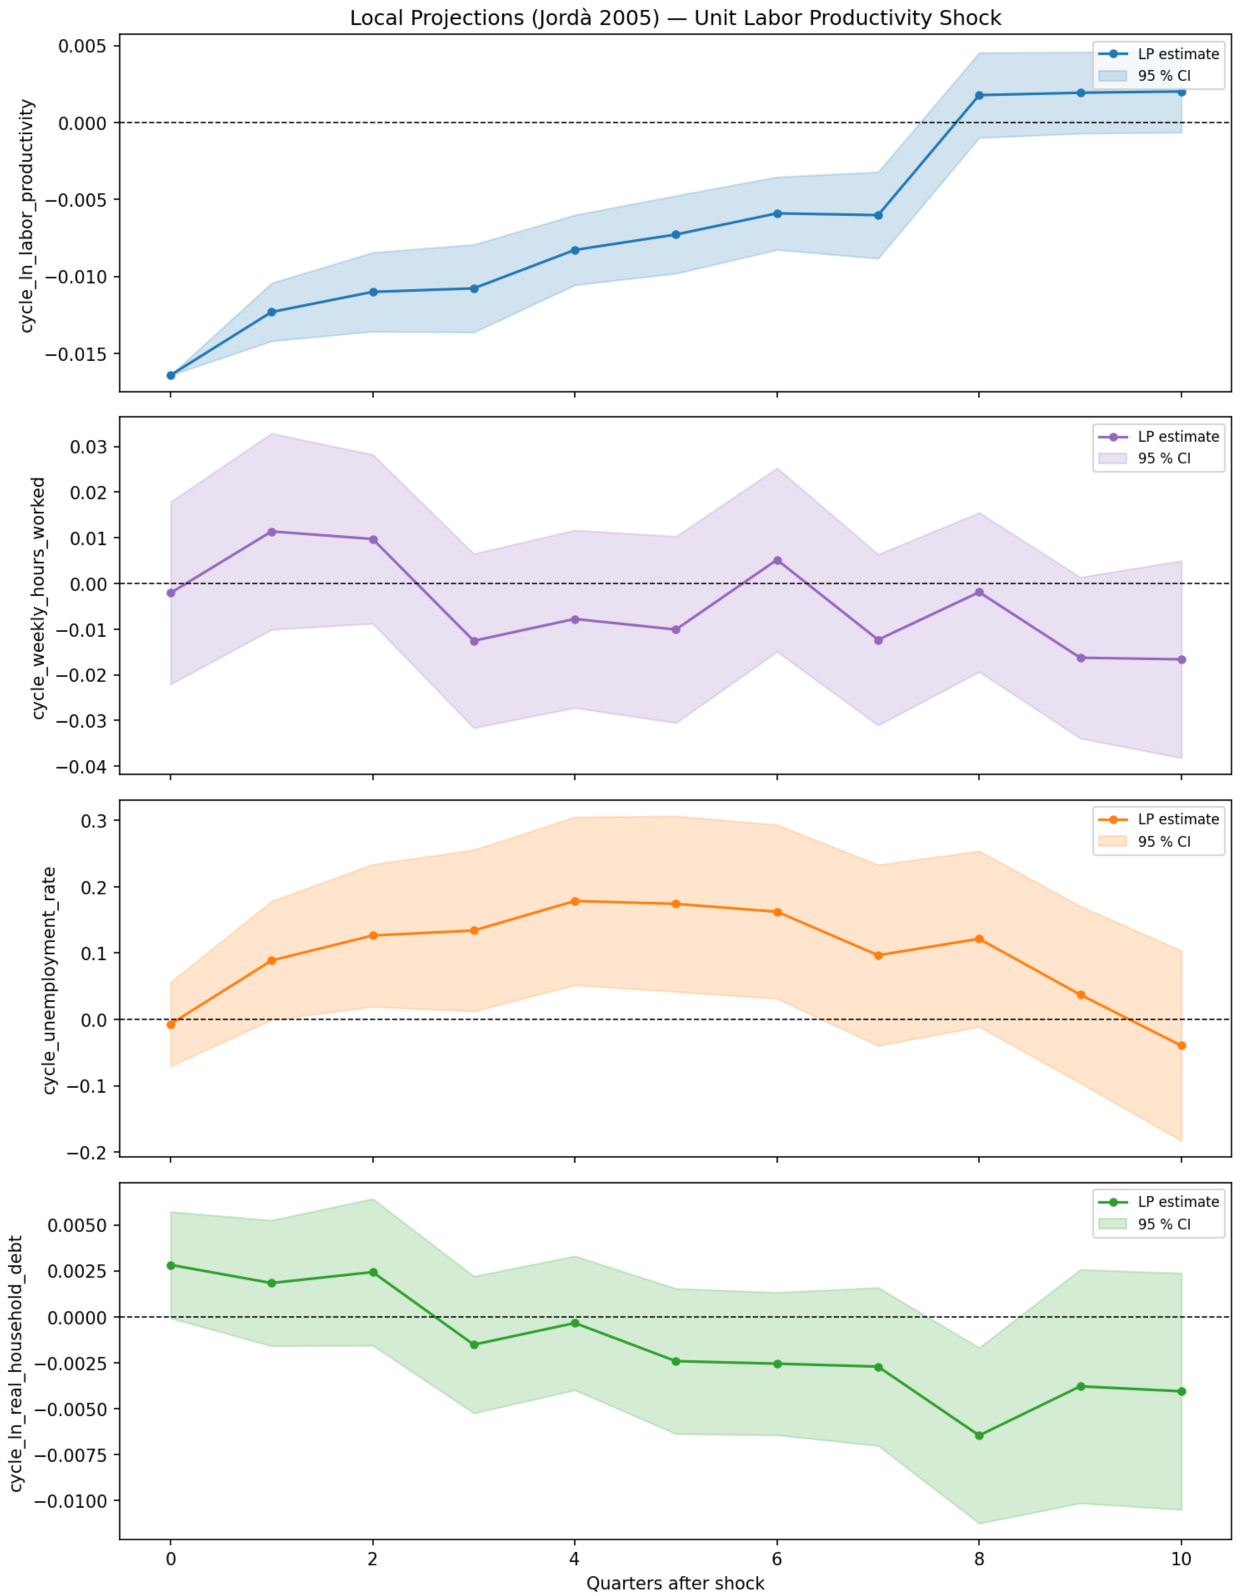

Figure 41: LP All Countries 2 Year filter

Figure (42) displays a VAR with the only difference being that the 4th variable is now the cyclical component of nominal household debt. Figure (43) displays a LP with the only difference being that the 4th variable is now the cyclical component of nominal household debt. In both these figures, there is little to no difference for all four variables. This was expected for the first three variables, as none of these were affected, and the fourth debt variable was somewhat unknown. The fact that the impulse response of nominal debt closely tracks that of real household debt suggests that productivity shocks operate through the extensive margin of employment, reducing labor demand, rather than via price effects. This distinction is more nuanced. (Gagliardone and Gertler 2023) shows that oil shocks can simultaneously affect employment and prices when oil is a complementary good for both household consumption and production. The fact that nominal debt responds as real debt implies that the labor productivity shock operates primarily through labor demand rather than through price changes. This is consistent with a productivity shock weakening credit channels via changes in employment status, rather than through inflationary price adjustments.

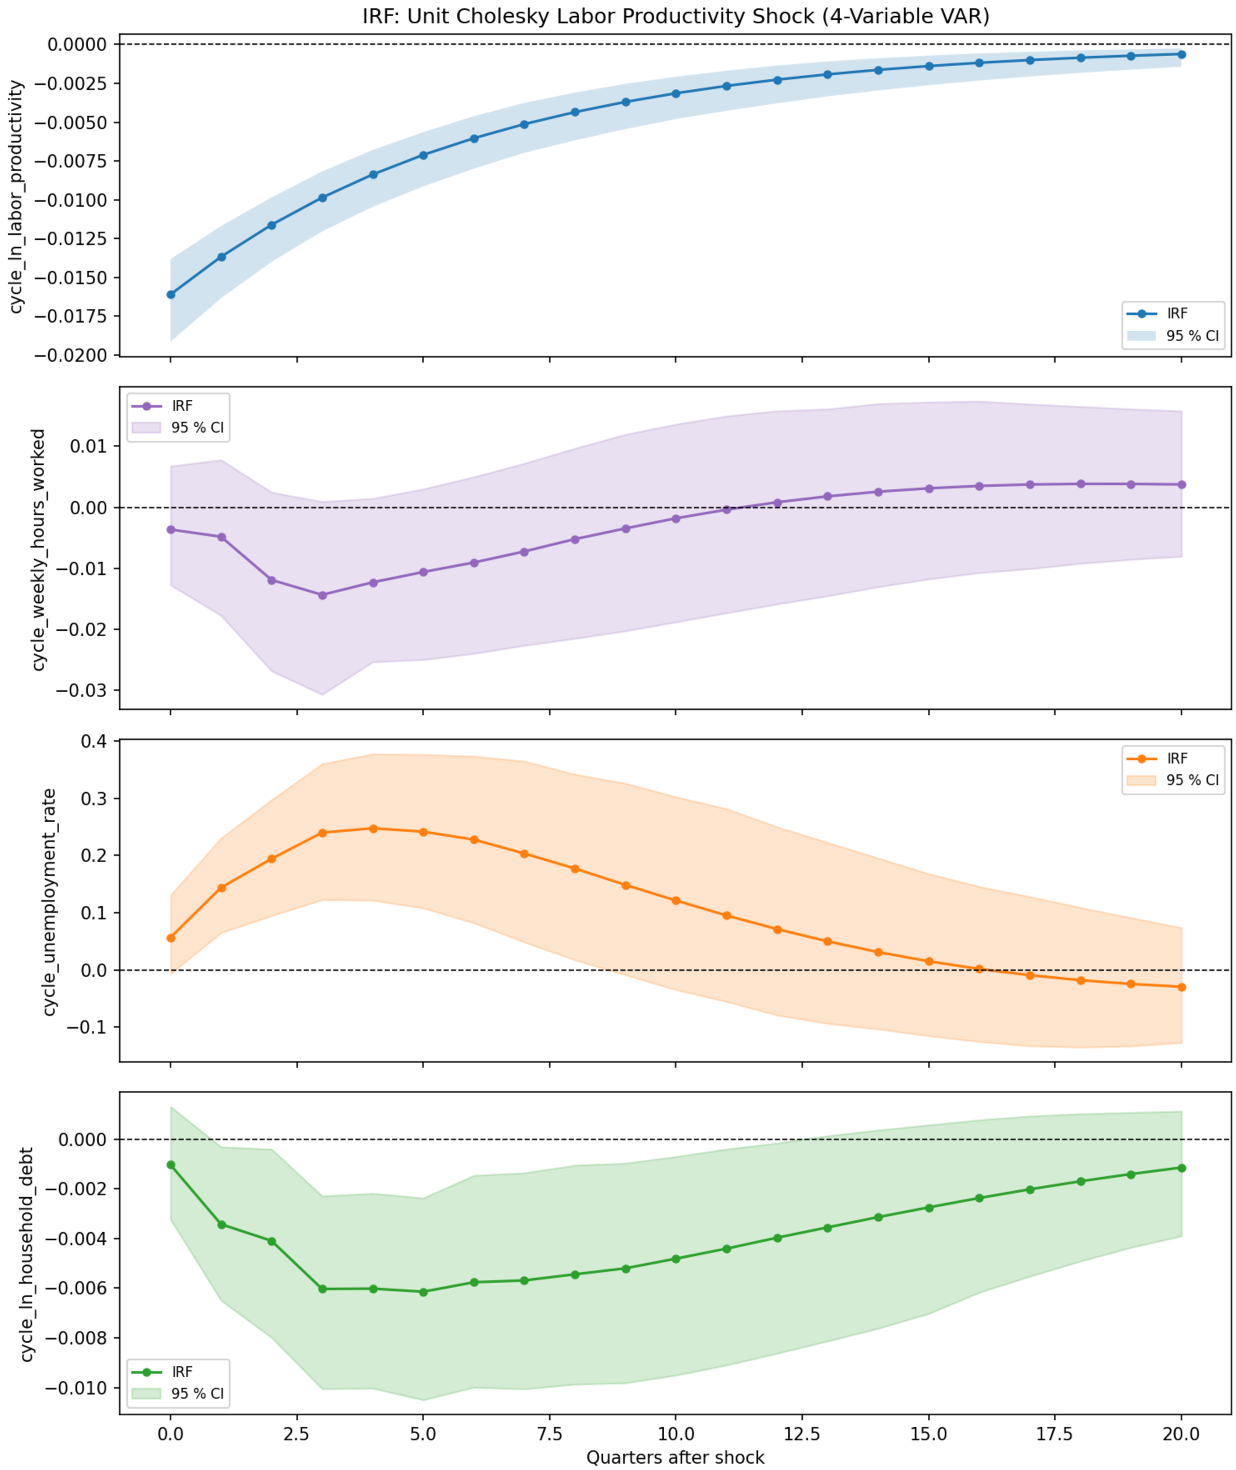

Figure 42: VAR All Countries Nominal Debt

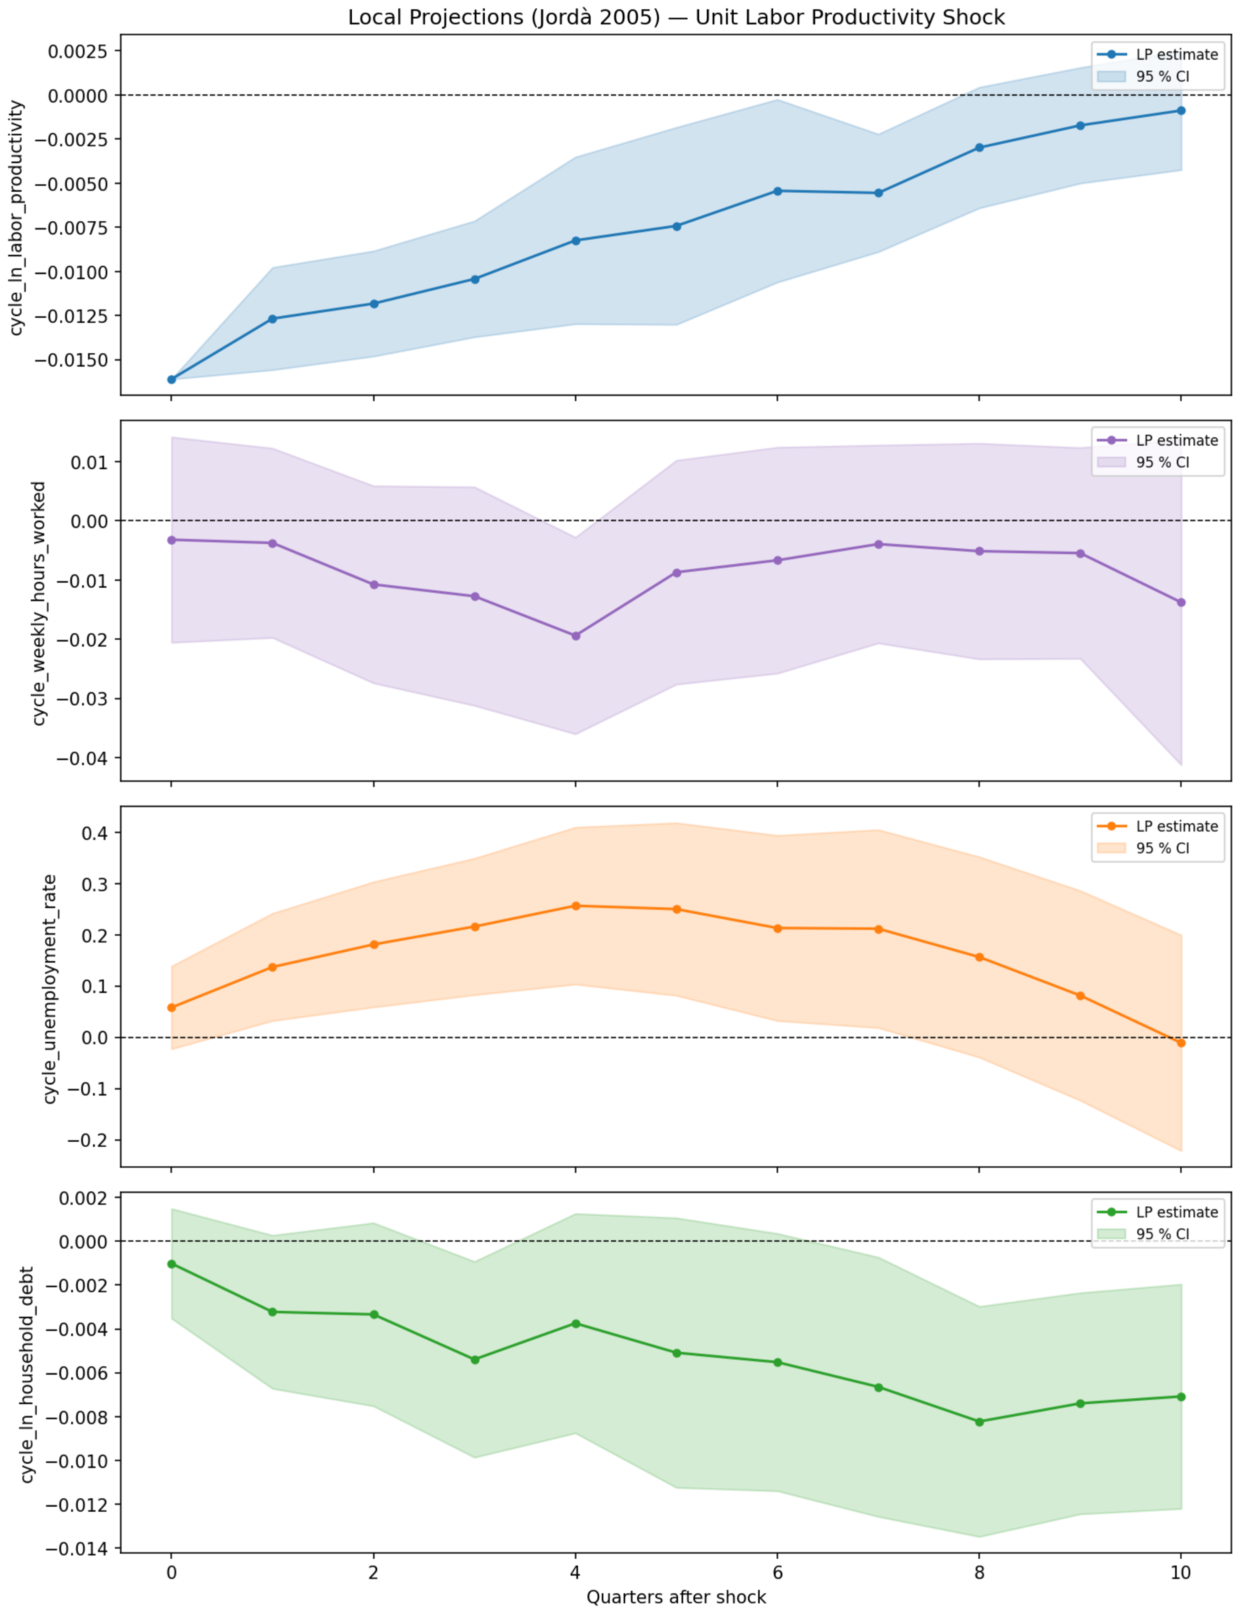

Figure 43: LP All Countries Nominal Debt

Figure (44) displays a VAR with the only difference being that the 4th variable is now the household debt to GDP. Figure (45) displays a LP with the only difference being that the 4th variable is now the household debt to GDP. The labor productivity, weekly hours worked, and unemployment show little to no difference, as expected. The difference is that the fourth variable, cycle household loans to GDP, increases right after the shock and decreases in both around 6 to 7 quarters after the shock. This makes sense, as the labor productivity shock reduces GDP more quickly than real household debt. This causes the initial increase in household debt to GDP, but as time goes on and labor productivity increases, so does GDP. This causes GDP to return to a value similar to the pre-shock level, and the decrease in household debt is evident in the nominal value of household debt, which is now roughly what it was at the start.

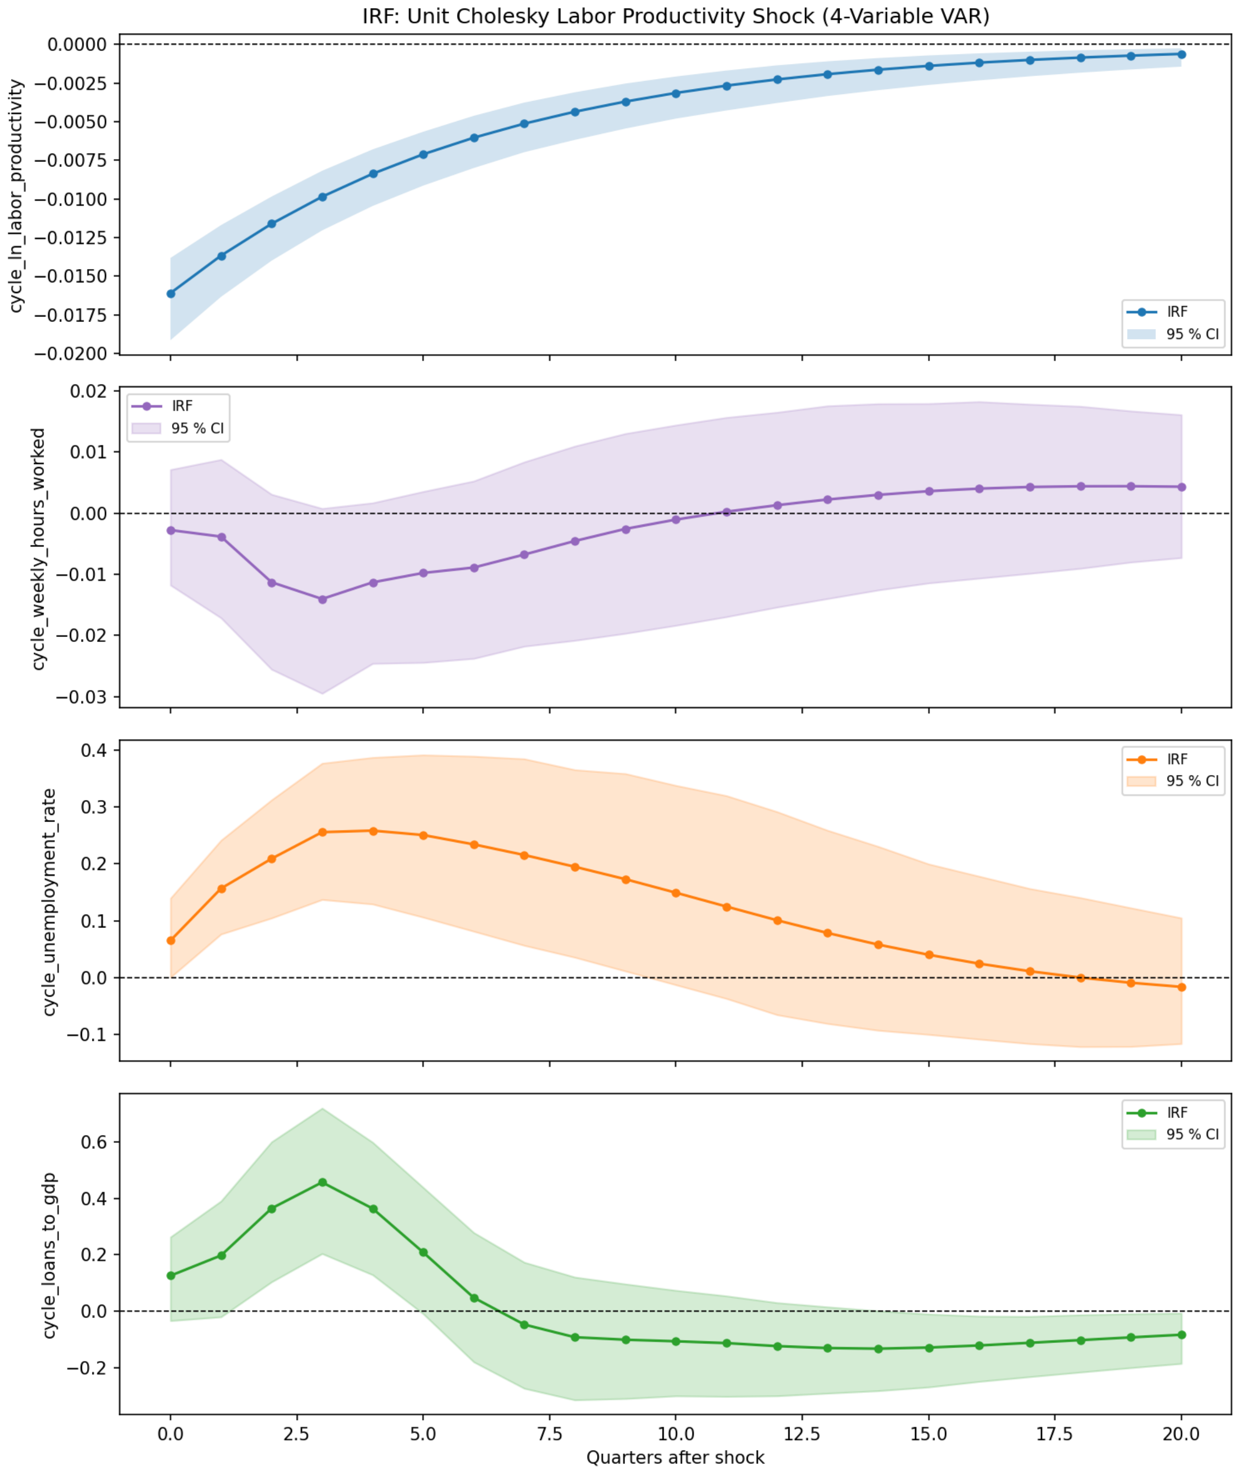

Figure 44: VAR All Countries Debt to GDP

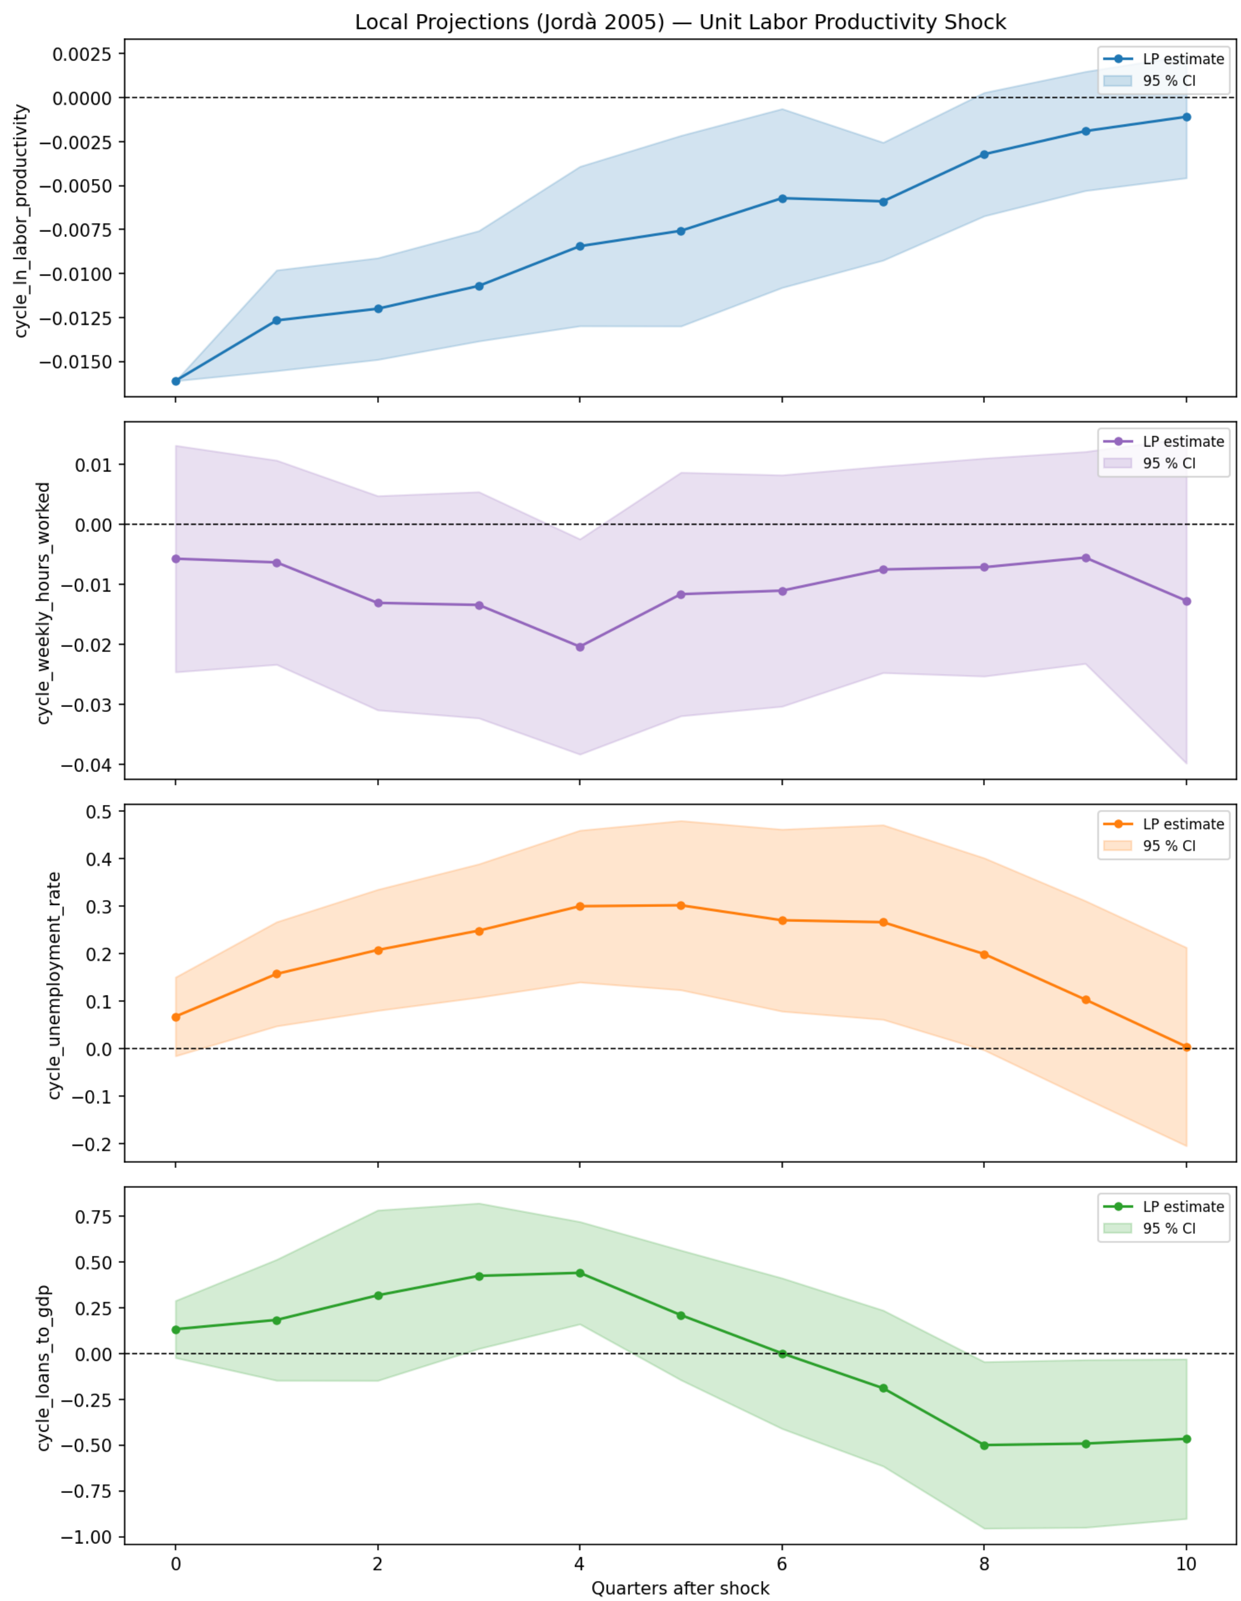

Figure 45: LP All Countries Debt to GDP