ARTICLE IN PRESS

Review of Economic Dynamics ••• (••••) •••–•••

Contents lists available at ScienceDirect

Review of Economic Dynamics

journal homepage: www.elsevier.com/locate/red

Demand, growth, and deleveraging [☆]

Brian Greaney [a,*], Conor Walsh [b]

[a] University of Washington, 410 Spokane Lane, Seattle, WA 98105, United States of America
[b] Columbia University, Kravis Hall, 665 W 130th St, New York, NY 10027, United States of America

ARTICLE INFO

Article history:
Received 27 April 2021
Received in revised form 22 August 2023
Available online xxxx

Keywords:
Deleveraging
Firm entry
Propagation
Growth

ABSTRACT

We present empirical evidence that weak household demand contributed to a reduction in firm entry in the Great Recession. Motivated by this evidence, we study the general equilibrium response of aggregate economic growth to a severe deleveraging event. To do so, we combine a standard incomplete markets model with a class of endogenous growth models. A large reduction in credit access causes the zero lower bound to bind, inducing a drop in demand via employment rationing. Decreased demand in turn lowers the return to entrepreneurship and innovation, endogenously lowering productivity. We find that a persistent recession induced by deleveraging can significantly influence growth in productivity. Our main result is a powerful feedback effect: households increase savings in response to future slow growth, exacerbate the fall in demand, and further slow the recovery.

1. Introduction

Deleveraging events can put significant downward pressure on demand and output. When access to credit tightens, households cut back on purchases and increase savings. This can induce a fall in labor demand, and with it aggregate production. An influential literature has argued that such a process played a key role in the employment losses of the 2008 financial crisis.

However, in many macroeconomic models of deleveraging, crises are brief. The aggregate economy adjusts very quickly to tighter credit limits, as consumers save rapidly to meet new debt thresholds. After a short period of pain lasting a few quarters, transitional demand pressures evaporate. In other words, internal propagation is absent. This stands in stark contrast to the Great Recession: unemployment rose to a post-war high in 2009, and took a decade to recover to its pre-crisis level.

In this paper, we consider a source of propagation that has received little attention, but in the data appears central to explaining employment losses. Creation of new firms fell sharply in the Great Recession, and also took years to recover. Given that new entrants contribute disproportionately to productivity growth and job creation, this had a severe impact on labor market recovery.

☆ We thank Giuseppe Moscarini for his continued guidance and support with this project. We also thank Costas Arkolakis, Sam Kortum, Michael Peters, Aleh Tsyvinski and Fabrizio Zilibotti for helpful comments, and Shirley Liu for help with the data. In addition, we thank the editor and an anonymous referee for suggestions which substantially improved the paper. Any opinions and conclusions expressed herein are those of the authors and do not necessarily represent the views of the U.S. Census Bureau. All results have been reviewed to ensure that no confidential information is disclosed.

* Corresponding author.
E-mail addresses: bg385@uw.edu (B. Greaney), conor.walsh.a@gmail.com (C. Walsh).

https://doi.org/10.1016/j.red.2023.08.004

We make two contributions. First, we show that declines in household demand contributed to a sharp reduction in firm entry during the Great Recession. We document that areas hit harder by house price declines in the 2008 recession saw greater falls in firm entry. Non-tradable sectors were much more affected than tradable sectors, suggesting that household demand was a key channel.

We then assess the strength of this propagation in general equilibrium. Coupling a standard model of household deleveraging with a model of endogenous growth and firm dynamics, we show quantitatively that reduced household spending in response to tighter credit limits is significantly amplified and prolonged by a fall in firm creation. A feedback effect occurs, where lower firm entry leads to lower growth, and households, expecting lower incomes in the future, cut back on current spending even further. Nominal rigidities translate this demand reduction into employment losses, further lowering the return to entrepreneurship and firm creation. A dynamic of this type can help explain some of the slow recovery in employment from the 2008 financial crisis.

Related Literature. An influential empirical literature has argued that household net worth shocks played a pivotal role in cross-sectional employment losses in the U.S. during the Great Recession (Campbell and Cocco, 2007; Mian and Sufi, 2011; Mian et al., 2013; Mian and Sufi, 2014; Kaplan et al., 2020). The greater response of non-tradable sectors to these shocks suggests a primary role for household demand, both through large propensities to consume out of housing wealth, and tightening household credit constraints. The literature has explored both of these channels, with the former being analyzed and modeled in Berger et al. (2018), Stroebel and Vavra (2019), and Kaplan et al. (2020).

Our quantitative work focuses on the latter. Influential work by Guerrieri and Lorenzoni (2017) computes transitional dynamics after a credit tightening in a Bewley-Huggett model. With flexible interest rates, they find that both the interest rate and output fall below their long-run values after the shock and then quickly adjust upward. Allowing agents to save in durable goods as well as bonds has little effect on these propagation patterns.

This lack of propagation is a common finding.[1] Justiniano et al. (2015) also use a two-type (patient and impatient) model. Household debt is in the form of mortgages on housing. They consider two deleveraging episodes: a loosening and subsequent tightening of the maximum loan-to-value ratio (which applies only to new loans), and an increase and subsequent decrease in preferences for housing (which effectively limits borrowing by changing collateral values while holding the collateral constraint constant). Neither episode generates a persistent decline in output.

We also speak to an active literature that has integrated growth and firm dynamics with models of the business cycle (Fatas, 2000; Barlevy, 2004; Comin and Gertler, 2006; Moran and Queralto, 2018; Anzoategui et al., 2019; Bianchi et al., 2019). Closest to this paper is work by Benigno and Fornaro (2018), who study how expectation-driven slumps can lead to persistently low productivity growth in a Schumpeterian setting. Our focus is studying explicitly how endogenously low growth can amplify and propagate employment losses caused by deleveraging, so we develop a quantitative model which can speak to the first-order features of household savings decisions, firm entry, and growth. On the empirical side, Adelino et al. (2015) and Davis and Haltiwanger (2019) study the effect of local house price changes on start-up collateral for new businesses. In our empirical work, we focus on the differential role of non-tradable vs. tradable sectors to disentangle the role of local demand.

2. Firm creation in the data

We first show that declines in household demand contributed to a sharp reduction in firm entry during the Great Recession. This reduction is known to be large. However, discussion is often bound up with the secular fact of declining entry rates since 1980 (Pugsley and Şahin, 2015), which has been linked to declining labor force growth.[2] To abstract from this, in Fig. 1a we plot the number of new establishments divided by the total labor force in each year. This measure of entry shows only a mild trend prior to the Crisis, but then suffers a decline of around 25% between 2007 and 2008. The size of the decline and its persistence has no parallel in the period for which we have data (late 1970s to present).

The collapse in entry and its failure to recover quickly has been costly. It is well-known that entrants contribute disproportionately to job creation and output growth (see e.g. Haltiwanger et al., 2013). We can directly trace the impact their post-Crisis absence has had in the data with a simple exercise. Using the publicly available data on firm creation from the Business Dynamics Statistics (BDS), we decompose overall employment growth since 1978 into three components: job creation due to entry of new firms, job destruction due to exit of existing firms, and net job creation of incumbent firms. These three components are plotted in Online Appendix A.1. This figure shows that job creation by entrants experienced a large decline below trend in 2008, and stayed depressed for an extended period. At the same time, job creation by incumbents contracted sharply, before quickly recovering to pre-Crisis levels.

In Fig. 1b, we construct a crude counterfactual scenario for employment. We hold job creation by entrants at its average from 2000 to 2006, and extrapolate forward from 2007 to 2016. We then add the actual series for net job creation by

[1] In other important work on deleveraging, lack of propagation is by design. Eggertsson and Krugman (2012) study the effects of deleveraging in an endowment economy with two types of agents: patient (lenders) and impatient (borrowers). In their model, the entire adjustment of the economy to a credit crunch occurs in one period. Work by Korinek and Simsek (2016) shares this feature, but endogenizes the buildup of debt before the deleveraging episode.

[2] See Karahan et al. (2018), Peters and Walsh (2022), and Hathaway and Litan (2014).

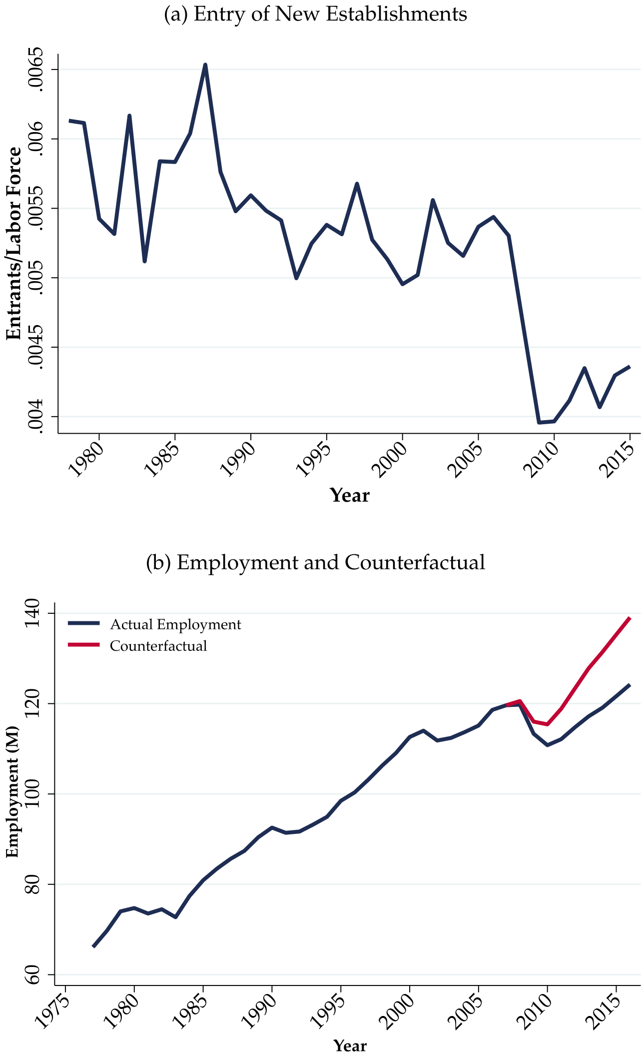

Fig. 1. Entry and The Great Recession. Note: This figure plots total new establishments in the United States by year, divided by an estimate of the total U.S. labor force in that year. The source data for establishments is the Business Dynamics Statistics database produced by the Center for Economic Studies at the U.S. Census Bureau. The data for labor force is from the Bureau of Labor Statistics. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

incumbents and job destruction due to exit. If job creation by entrants had held at its Pre-Crisis levels, employment would have been higher by some 15 million jobs by 2016.

The decline in firm entry also had a strong spatial component. In particular, counties and Metropolitan Statistical Areas (MSAs) in the U.S. that suffered larger house price falls saw larger falls in firm entry. In Fig. 2, we show the number of operating establishments relative to 2005 across two groups of counties, split by whether that county experienced a house price decline above or below the median between 2007 and 2009. While establishment counts were growing similarly in the years before the Crisis, counties that were particularly hard hit by house price declines saw larger falls in establishment counts after the Crisis. Moreover, establishment counts stayed depressed for longer in these counties, mirroring the dynamics of employment uncovered in Mian and Sufi (2014).

This was mostly driven by a decrease in establishment creation. In Table 1, we employ the public data on establishment creation provided by the Census Bureau in its BDS product, which provides counts of new establishments across MSAs. We regress the percentage change in the number of new establishments from 2007 to 2010 on the percentage change in house prices in that MSA during that period. The elasticity is large. Going from the highest quartile of house price change to the lowest sees establishment creation decline by 6%. In the second column of Table 1, we instrument house price falls with the local housing supply elasticity from Saiz (2010), following the strategy pioneered by Mian and Sufi (2014). The coefficient is virtually unchanged, suggesting housing price declines led to substantial reductions in local establishment creation.

In the second panel of Table 1, we examine job creation by entrants. This coefficient is very similar to that observed for raw establishment creation, suggesting most of the effect on jobs came on the extensive margin. As such, loss of new

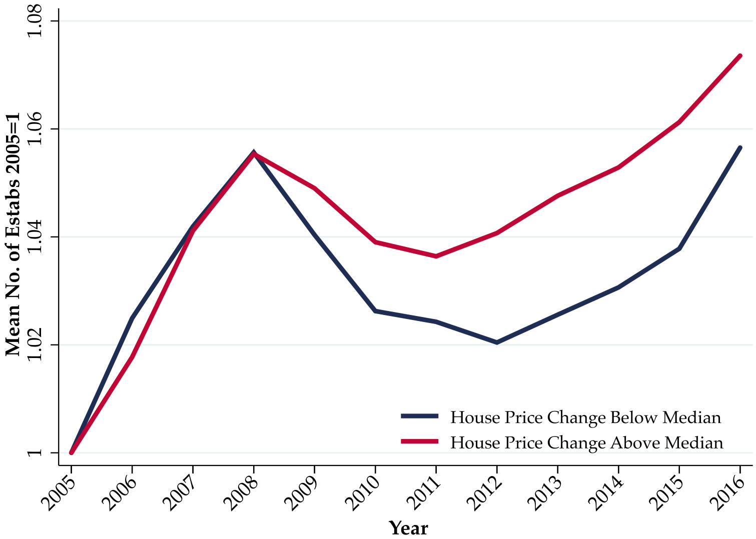

Fig. 2. Establishment Counts Across Counties. Note: This figure plots mean total establishments within a county relative to 2005. The mean is taken across all counties of the U.S. split into two groups, based on whether that county had a mean house price change above or below the median between 2008 and 2009. The source data for establishment counts is the Quarterly Census of Employment and Wages (QCEW) published by the Bureau of Labor Statistics. The data for house prices at the county level comes from Zillow.

Table 1
Change in House Prices and Establishment Entry at the MSA Level, 2006-2010.

	Dependent variable:			
	% Change in Entry of Establishments		% Change in Job Creation by Entrants	
	OLS	*IV*	*OLS*	*IV*
% Change in Avg. House Price	0.145***	0.174***	0.138**	0.183
	(0.023)	(0.055)	(0.055)	(0.130)
Constant	−0.252***	−0.249***	−0.265***	−0.261***
	(0.004)	(0.007)	(0.010)	(0.016)
Observations	239	239	239	239
R^2	0.137	0.131	0.022	0.019

Note: Measures of establishment entry and entrant job creation come from the Business Dynamics Statistics. House prices at the MSA level come from the FHFA All Transactions Index. Housing supply elasticity for the IV comes from Saiz (2010). Robust standard errors in parentheses. *** significant at the 1% level, ** significant at the 5% level, * significant at the 10% level.

establishments fed directly into reduced job creation, with no countervailing expansion of employment by those who did enter.

There are two primary economic mechanisms that could explain these results. The first is what we focus on in this paper: demand reductions driven by consumer deleveraging and wealth effects. Mian et al. (2013) document large elasticities of consumption with respect to housing wealth during this period. All else equal, lower consumer spending should lead to reduced expected profits for new entrants, and reduce the incentive to start a business. Through their effect on consumer spending, house price falls might directly impact firm profitability. An alternative explanation is that lower house prices reduced the value of collateral available to small investors, tightening financial constraints on entrants.

To distinguish these stories, we employ the Census micro-data in the Longitudinal Business Database (LBD). The LBD contains information on every employer establishment in the United States since 1978, including age, employment, location and industry. We use this information to construct detailed measures of entry of new establishments at the county level. In particular, we classify industries as either Non-Tradable, Tradable, or Other, using the 4-digit NAICS categorization of Mian and Sufi (2014) based on geographic concentration. If demand for a firm's products is an important determinant of the decision to enter, we would expect Non-Tradable industries to suffer proportionally more when house prices decline, since local consumer demand is important for firms in these industries.

We test this by regressing the percentage change in establishment entry between 2006 and 2010 on two measures of house price changes. The first is a county-level measure of net worth from Mian and Sufi (2014), and the second is average house prices at the county level from Zillow. Panel A in Table 2 shows the results for Non-Tradable industries. Counties that saw larger house price declines saw establishment entry in Non-Tradables decline by more. In Columns II and V, we

Table 2
House Prices and Change in Establishment Entry at the County Level, 2006-2010.

	% Change in Establishment Entry, 2006-2010					
Panel A: Non-Tradable Entry	I	II	III	IV	V	VI
% Change in Average Net Worth	0.307***	0.364***	0.365***			
	(0.078)	(0.109)	(0.129)			
% Change in Average House Price				0.126**	0.130*	0.116
				(0.055)	(0.067)	(0.090)
Industry Employment Controls	No	Yes	Yes	No	Yes	Yes
County Controls	No	No	Yes	No	No	Yes
R-squared	0.005	0.044	0.072	0.010	0.046	0.068
Observations	850	850	850	950	950	950
Panel B: Tradable Entry						
% Change in Average Net Worth	-0.302	-0.340	-0.218			
	(0.449)	(0.177)	(0.202)			
% Change in Average House Price				-0.148	-0.163	-0.163
				(0.162)	(0.511)	(0.090)
Industry Employment Controls	No	Yes	Yes	No	Yes	Yes
County Controls	No	No	Yes	No	No	Yes
R-squared	0.002	0.016	0.039	0.001	0.015	0.040
Observations	850	850	850	950	950	950

Note: Definitions of Tradable and Non-Tradable industries are taken from Mian and Sufi (2014), using industry-level geographic concentration. House price data is from Zillow. Each N is rounded according to Census Bureau rules on disclosure avoidance. County level controls include percent black, average age, mean wealth, the pre-Crisis unemployment level and an indicator for an urban area. Robust standard errors in parentheses. *** significant at the 1% level, ** significant at the 5% level, * significant at the 10% level.

include controls for pre-Crisis industry employment shares at the county level, while in Columns III and VI we also include county-level demographic controls. Doing so does not change this conclusion.

In contrast, the results for Tradable Entry in Panel B have the wrong sign, and are not statistically significant in any of the six estimated models. It appears that entry during this period for Tradable industries was unaffected by declines in local house prices. This suggests that credit constraints on entrepreneurs due to declines in housing equity are not the driving force behind our results.

We consider a number of alternative specifications in the Online Appendix. In Table 4 we consider the net change in the number of establishments (i.e. new entrants minus exiters). In Table 5 we restrict the analysis to new firms only (i.e. establishments that do not belong to an incumbent firm). In Table 6 we consider the percentage change in job creation by new entrant establishments. All these alternative specifications yield similar results.

Together, these results suggest an important role for demand and spending in driving firm and establishment creation. In both theory and data, firm creation is a key source of growth. Young firms generate rapid expansion in employment and productivity upgrading, and their absence in a recession can significantly influence macroeconomic dynamics. We now turn to theory to examine this link.

3. Model

In this section, we develop a theoretical framework for analyzing the interrelationship between household demand and productivity growth. The model can broadly be categorized into two modules: a firm module which determines productivity growth, and a household module which determines aggregate demand. Combining them into a unified setting allows us to quantify the importance of the mechanisms we have outlined above in general equilibrium.

3.1. Firms

The firm side of the model uses and extends the framework developed in Atkeson and Burstein (2019), which nests both early Schumpeterian growth models and expanding varieties models. This framework is quite flexible, and can easily be mapped to settings which consider multi-product firms, as in Klette and Kortum (2004), Lentz and Mortensen (2008) and Garcia-Macia et al. (2019). These models are widely used to study the sources of productivity growth with firm-level data, and have become the workhorse of the growth literature.

3.1.1. Production

Time t is continuous. There is a single, final good which serves as the numeraire. This final good is composed of differentiated intermediates, which are produced by monopolistically competitive firms. Intermediates are aggregated into the final good by perfectly competitive firms using the technology

$$Y_t = \left(\int_0^{M_t} q_t(i)^{\frac{\sigma-1}{\sigma}} di \right)^{\frac{\sigma}{\sigma-1}}, \tag{1}$$

where M_t is the mass of firms at time t, $q_t(i)$ is the quantity of firm i's intermediate. Each intermediate variety is characterized by the frontier technology $z \in \mathbb{R}_+$ which is available to produce it. The production technology is $q(z) = z^{\frac{1}{\sigma-1}} l$, where l is labor. Labor is supplied by households in a competitive labor market at wage w_t. Aggregate labor supply is denoted by L_t, and is determined in general equilibrium.

In this setting, markups are constant across firms in all periods and are equal to $\overline{\sigma} \equiv \sigma/(\sigma-1)$. This implies that our economy admits a simple aggregate production function representation of

$$Y_t = (M_t A_t)^{\frac{1}{\sigma-1}} L_t^P = \overline{\sigma} w_t L_t^P, \tag{2}$$

where $A_t \equiv \int_{\mathbb{R}_+} z m_t(z) dz$ is average productivity and L_t^P is labor devoted to production of intermediates.

In our baseline model, each firm owns only a single variety. An endogenously determined flow M_t^e of new firms enter at t. With probability $\delta_e \in [0,1)$, they receive a proportional improvement Δ on the technology for producing an existing firm's product. With probability $1-\delta_e$, they invent a good that is wholly new to society. The productivity of a firm that invents a new good is ωA_t, where ω is drawn from a time-invariant distribution $\Gamma(\omega)$ with mean $\bar{\omega}$. This introduces a spillover from current operating technology to the draws of current entrants, and, as the economy grows, ensures that the level of technology of entrants improves alongside the aggregate economy.

An incumbent firm that does not have its product innovated upon receives stochastic improvements to its productivity. We assume that with Poisson rate ϕ, the firm's productivity increases by the proportional factor Δ. In the baseline model, this rate is exogenous and costless.[3] Lastly, we assume that products die with an exogenous Poisson rate δ_0. The knowledge of how to produce the product is then lost to society.

Given these assumptions, it can be shown that average productivity evolves according to

$$g_t^a \equiv \frac{\dot{A}_t}{A_t} = (\Phi_1 + m_t^e \Phi_2), \tag{3}$$

where $m_t^e = M_t^e / M_t$, and

$$\Phi_1 = \phi(\Delta - 1),$$
$$\Phi_2 = \delta_e(\Delta - 1) + (1-\delta_e)(\bar{\omega} - 1).$$

As this expression makes plain, the growth in average productivity is driven by the improvements that incumbents make, captured in Φ_1, and the contribution from entry, captured in Φ_2. The relative contribution of entrants will depend on how productive they are relative to the average incumbent, which is captured by the weighted average of the Schumpeterian step size Δ, and the mean productivity of an entirely new product, $\bar{\omega}$.

3.1.2. Firm ownership

We assume that firms are owned by a representative capitalist. The capitalist uses the firms' profits for consumption and to finance the creation of new firms. This capitalist cannot work, and has preferences only over consumption C_t^K, which it discounts at rate ρ_k, given by

$$U^K = \int_0^{\infty} e^{-\rho_k t} \log(C_t^K) dt. \tag{4}$$

The capitalist is fully-forward looking, and invests in its portfolio of firms to maximize discounted lifetime utility. We assume it does so competitively, as would a mass of identical capitalists, and does not internalize the effects of creative destruction and increased competition that results from new investment.

3.1.3. Firm value function

We now turn to characterizing the value of each firm. As is standard for these models,[4] period profits of firms with productivity z at time t are a constant fraction of revenues:

[3] In Online Appendix C.1, we endogenize these movements by considering purposeful investment by incumbents in improving their productivity.

[4] The CES production function for the final good in (1) implies that all firms would like to charge a markup over marginal cost of $\overline{\sigma}$ if their pricing were unconstrained. However, as in all models with a Schumpeterian element, the question of what happens to the incumbent who experiences creative destruction means that pricing may not be unconstrained; the new firm may have to consider the previous leader in the good when setting prices. For this section, if an entrant firm innovates on top of another firm's product, we assume that the improvement is sufficiently drastic that the previous leader's marginal cost is above the new firm's price at their optimal markup. The restriction on parameters for this to be the case is $\Delta \geq \overline{\sigma}$. Heterogeneous markups with limit pricing as in Peters and Walsh (2022) can easily be handled, and is discussed briefly in Online Appendix C.2.

$$\pi_t(z) = \frac{1}{\sigma}\frac{Y_t}{A_t M_t} z. \tag{5}$$

This is a well-known expression, but in our context it deserves some discussion. First, all else equal, increases in either average productivity A_t or the total mass of firms M_t decrease profits. This effect operates by bidding up the real wage, increasing an individual firm's marginal cost for a given z. In this way, aggregate growth erodes the cost advantage of a firm for a given z. Second, increases in aggregate production Y_t scale firm profits one-for-one. In equilibrium, aggregate production must be equal to the aggregate demand for the final good. As such, aggregate demand is a key determinant of the profitability of a firm.

We denote the value to the capitalist of a given firm with productivity z by $V_t(z)$. This value satisfies the Hamilton-Jacobi-Bellman (HJB) equation

$$r_t^k V_t(z) = \pi_t(z) + \underbrace{\dot{V}_t(z)}_{\text{Capital gains}} + \underbrace{\phi(V_t(\Delta z) - V_t(z))}_{\text{Innovation}} - \underbrace{(\delta_e m_t^e + \delta_0)V_t(z)}_{\text{Death and creative destruction}}. \tag{6}$$

Here r_t^k is the return on the optimal portfolio of the capitalist, who can perfectly diversify idiosyncratic firm risk since each firm is infinitesimal.

3.1.4. Entry and growth

We assume that entry of new products requires hiring $\bar{\epsilon}$ units of labor for research.[5] As such, labor market clearing requires

$$L_t^p + \bar{\varepsilon} M_t^e = L_t. \tag{7}$$

The cost and value of firm entry are related via the free entry condition

$$w_t\bar{\epsilon} \geq (1-\delta_e)\int_{\mathbb{R}_+} V_t(z)dG_t(z) + \delta_e \int_{\mathbb{R}_+} V_t(\Delta z)m_t(z)dz, \tag{8}$$

where $G_t(z) = \Gamma(\frac{z}{A_t})$ is the distribution of productivity when creating a new product, as detailed above. (8) holds with equality if and only if firm entry M_t^e is strictly positive. The first term on the right hand reflects the possibility that the entrant invents a wholly new product, while the second term reflects the possibility that it innovates on top of an existing product. The distribution of productivity at time t, here captured through the density $m_t(z)$, fully determines the productivity of the entrant in the case they receive an innovation on an existing product.

3.2. Households

The household side of the model is closely related to the standard incomplete-markets framework with idiosyncratic labor income risk and a borrowing limit. The main novelty is that there is aggregate growth. As originally noted by Aiyagari (1994), detrending by a constant aggregate growth rate leaves the distribution of assets, consumption, and income stationary on the balanced growth path. Outside of the balanced growth path, households must forecast the full path for wage growth in order to make their decisions, which is taken into account in the definition of equilibrium.

In line with the empirical evidence shown previously, we focus on deleveraging as the source of aggregate demand shortfalls. An unexpected tightening of the borrowing limit causes households to increase their savings, and, all else equal, lowers the equilibrium interest rate on debt. In the presence of a zero lower bound, or more generally an inflexible interest rate, this can create a shortfall of demand. In particular, with nominal rigidities the real rate is not fully flexible (see e.g. Korinek and Simsek, 2016 and Auclert and Rognlie, 2018), and cannot adjust downwards to offset the shortfall in aggregate demand induced by deleveraging. As such, the bond and goods markets will not clear without an adjustment in another part of the economy. We assume this adjustment occurs through classical unemployment. The amount of labor employed is what is demanded for production at the given real interest rate. To deal with the excess supply, employment is rationed among households, and their real spending on consumption falls. This causes a fall in firm profits, and a corresponding reduction in entry and product creation.

While undoubtedly special, such a setting allows us to fully specify the source of reductions in demand, and consider in general equilibrium any feedbacks from fluctuations in the growth rate back on to demand. A natural amplification channel occurs in this model, whereby consumers anticipate the coming reduction in growth, and increase their savings. With an inflexible real rate, this only worsens the slump in employment, since demand for consumption falls further.

[5] One could, alternatively, specify that a cost in units of the final good must be spent to begin a new firm, or that a research good that uses both labor and units of the final good is needed, as in Atkeson and Burstein (2010). Both of these formulations can be made isomorphic to our current setup.

3.2.1. Environment

There is a unit mass of infinitely-lived households, and no population growth. Households discount the future at rate ρ_h and have flow utility function

$$u^h(C) = \frac{C^{1-\gamma}}{1-\gamma},$$

where C is consumption. Households supply labor inelastically for the competitive wage w_t, but are subject to stochastic shocks to their labor endowment. Individual labor endowments ℓ follow a continuous-time Markov process with states $\{\ell_1, ..., \ell_J\}$ and switching intensities $\lambda_{jj'}$. Households are only allowed to supply a fraction L_t of their labor endowment. L_t is nonstochastic and may be less than 1 if there is labor rationing. They can borrow and save via a risk-free bond B, subject to an ad-hoc borrowing constraint $B \geq \underline{B}_t$. The (real) return on bonds is denoted by r_t^h. Bonds are in zero net supply.

Households cannot hold equity in firms, which are fully owned by the representative capitalist. This separates somewhat the interest rate on households' savings r_t^h, which governs their consumption and saving decisions, and the interest rate on equity r_t^k, which governs growth and the investment decisions of firms. This assumption, while somewhat stark, is empirically relevant. In particular, Wolff (2017) finds that the 93% of all stocks and 94% of business equity is held by 10% of households, which makes a limited participation model an appealingly parsimonious description of the data.[6]

3.2.2. Nominal rigidity

We assume the presence of two frictions in this economy, which together imply that a deleveraging shock can have real effects. First, there exists a lower bound on the nominal interest rate i_t, motivated by the presence of some background quantity of cash which households would prefer to hold in the presence of a negative nominal interest rate. By itself, this bound has no implications for real allocations if prices are perfectly flexible, since the real interest rate is what determines allocations, and with flexible price setting inflation will perfectly adjust to ensure full employment (see Krugman et al., 1998 and Korinek and Simsek, 2016). It is only when combined with nominal rigidities that the zero lower bound becomes relevant. In such a world, the bound on the nominal rate causes a bound on the real rate.

The nominal rigidity we assume is that nominal wages w_t^n are growing at the BGP rate of g at all times, such that $w_t^n = w_0 e^{gt}$. While undoubtedly extreme, this simple form of rigidity allows us to analyze a global solution for equilibrium, where popular alternatives such as staggered price setting by firms do not.[7] With firms free to set any prices they like, the nominal rigidity in wages will characterize the dynamics of the price level. The aggregate price index P_t is the usual CES index incorporating the payment of marginal costs in nominal wages:

$$P_t = \frac{w_t^n}{\overline{\sigma}(M_t A_t)^{\frac{1}{\sigma-1}}}.$$

This implies that inflation is given by $\pi_t = g - g_t$. But then, since the real rate is equal to the nominal rate minus inflation, the zero lower bound corresponds to a bound on the real rate of

$$r_t^h \geq g_t - g.$$

So what if the real rate hits this bound at any given time? In that case, the asset and goods markets will not clear. We assume that if this happens, the economy adjusts by creating involuntary unemployment. In particular, when the zero lower bound binds, L_t adjusts so that the bond and goods markets clear. The zero lower bound constraint can be summarized by

$$\begin{aligned} & r_t^h > g_t - g \quad \text{and } L_t = 1 \\ \text{or } & r_t^h = g_t - g \quad \text{and } L_t < 1 \end{aligned} \tag{9}$$

3.2.3. The household's value function

Let $V_t^h(B, \ell)$ denote the value function of a household with wealth B and labor endowment ℓ. In Online Appendix E.1, we show that the household's value function satisfies the HJB equation

$$\rho_h V_t^h(B, \ell_j) = \max_C \frac{C^{1-\gamma}}{1-\gamma} + \partial_B V_t^h(B, \ell_j)(w_t^n L_t \ell_j + r_t^h B - C)$$

[6] Guvenen (2006, 2009) shows that such a structure can reconcile key micro evidence on the low elasticity of intertemporal substitution of the majority of households with the relatively high value needed to match growth and output fluctuations in standard frameworks.

[7] A possible alternative to study the real effects of deleveraging would be to employ the Heterogeneous Agent New Keynesian model (as in Kaplan et al. (2018)). There, nominal rigidities in prices also divorce the real interest rate from the nominal interest rate, and allow the central bank to influence output through the consumers' Euler equations, thereby affecting labor demand and aggregate output. However, outside of a simple expanding variety model, as explored in Moran and Queralto (2018), pricing decisions a la Calvo or Rotemberg become intractable when combined with exogenous or endogenous improvements in productivity post-entry, as occurs in our class of models.

$$+\sum_{j'=1}^{J}\lambda_{jj'}[V_t^h(B,\ell_{j'})-V_t^h(B,\ell_j)]+\dot{V}_t^h(B,\ell_j)$$

$$\text{subject to } w_t^n L_t\ell_j + r_t^h B - C \geq \dot{\underline{B}}_t \text{ if } B=\underline{B}_t$$

In order to obtain a stationary problem, define the detrended variables $b_t \equiv B_t/w_t^n$, $c_t \equiv C_t/w_t^n$, and $v_t^h(b,\ell_j) \equiv v_t^h(b,\ell_j)/(w_t^n)^{1-\gamma}$. As shown in Online Appendix E.2, the detrended HJB can then be written as

$$[\rho_h-(1-\gamma)g_t]v_t^h(b,\ell_j)=\max_c \frac{c^{1-\gamma}}{1-\gamma}+\partial_b v_t^h(b,\ell_j)[L_t\ell_j+(r_t^h-g_t)b-c] \tag{10}$$
$$+\sum_{j'=1}^{J}\lambda_{jj'}[v_t^h(b,\ell_{j'})-v_t^h(b,\ell_j)]+\dot{v}_t^h(b,\ell_j)$$
$$\text{subject to } L_t\ell_j+(r_t^h-g_t)b-c \geq \dot{\underline{b}}_t \text{ if } b=\underline{b}_t$$

The household's policy function $c_t(b,\ell_j)$ determined by (10) governs the evolution of aggregate bond-holdings in the economy. Let $\varphi_t(b,\ell_j)$ denote the unconditional joint density of wealth and labor endowments. In Online Appendix E.3, we show that this density satisfies the Kolmogorov forward (KF) equation

$$\dot{\varphi}_t(b,\ell_j)=-\partial_b\{\varphi_t(b,\ell_j)[L_t\ell_j+(r_t^h-g_t)b-c_t(b,\ell_j)]\} \tag{11}$$
$$-\left(\sum_{j'=1}^{J}\lambda_{jj'}\right)\varphi_t(b,\ell_j)+\sum_{j'=1}^{J}\lambda_{j'j}\varphi_t(b,\ell_{j'})$$

Since bonds are in zero net supply, the bond market clearing condition is

$$0=\sum_{j=1}^{J}\int_{\underline{b}_t}^{\infty} b\varphi_t(b,\ell_j)db \tag{12}$$

3.3. Equilibrium

We now define the equilibrium.

Definition 1. *An equilibrium is a path for interest rates on equity r_t^k and bonds r_t^h, wages w_t, production labor L_t^p, aggregate labor L_t, output Y_t, mass of firms M_t, capitalist policy function C_t^K, firm value functions $V_t(z)$, household policy functions $c_t(b,\ell)$, value functions $v_t^h(b,\ell)$, and densities of state variables $\varphi_t(b,\ell)$, such that:*

1. *The output market clearing condition (2) holds*
2. *The firm value function satisfies the HJB equation (6)*
3. *Capitalist consumption decisions maximize (4) subject to the labor market allocation condition (7)*
4. *The free entry condition (8) holds*
5. *The zero lower bound constraint (9) holds*
6. *The household's policy and value functions satisfy the HJB equation (10)*
7. *The density of household state variables satisfies the KF equation (11)*
8. *The bond market clearing condition (12) holds*

A balanced growth path (BGP) equilibrium is an equilibrium in which all equilibrium objects and distributions are time-invariant. The numerical algorithm we use to compute the off-BGP dynamics is described in the Online Appendix.

3.4. Firm dynamics, demand and the aggregate growth rate

The model admits a unique balanced growth path in which demand and market size plays no role, due to the absence of scale effects on the growth side of the model (see a discussion in Peters and Walsh, 2022). Both the entry rate and the aggregate growth rate are independent of total demand. It is only off the balanced growth path that sudden falls in profitability can cause a decline in entry.

To see this, we begin with a result that allows for straightforward characterization of the model both in and out of the balanced growth path. Proofs of the results of this section are shown in Online Appendix B.

Lemma 1. *The value of a firm is linear in z at all times: $V_t(z)=B_t z$ for some scalar $B_t>0$.*

When combined with the free entry condition, this fact is powerful. It implies that the value of the firm for all levels of productivity is proportional to the entry cost. This allows us to characterize how movements in current profitability impact macroeconomic aggregates in a tractable manner. We now consider how the growth rate is determined off the BGP. A key ratio that governs aggregate movements in productivity is the ratio of current profits to firm value:

$$\Theta_t(z) \equiv \frac{\pi_t(z)}{V_t(z)}.$$

Given Lemma 1 and the fact that profits are linear in z (equation (5)), this ratio is independent of z: $\Theta_t(z) = \Theta_t$. As such, we can evaluate it at mean entrant productivity $\bar{\omega}A_t$, which yields

$$\Theta_t = \frac{Y_t\bar{\omega}}{\sigma M_t V_t(\bar{\omega}A_t)}. \tag{13}$$

Now, with some manipulation of the free entry condition in (8) along with Lemma 1, it can be shown that, as long as firm entry is strictly positive, the value of the firm at mean entrant productivity is pinned down by the entry cost, such that

$$V_t(\bar{\omega}A_t) = \frac{w_t\bar{\epsilon}}{1-\delta_e+\delta_e\Delta/\overline{\omega}}.$$

Furthermore, our choice of the final good as numeraire implies that $w_t = Z_t$, so that we arrive at

$$\Theta_t = \frac{Y_t\overline{\omega}}{\sigma\tilde{\varepsilon}M_t Z_t},$$

where $\tilde{\epsilon} \equiv \bar{\epsilon}/(1-\delta_e+\delta_e\Delta/\overline{\omega})$. This expresses a balance of demand for the firm's good, coming from aggregate spending, and competition from other products. The free entry condition ensures that at all times, the detrended value of the firm in the denominator of (13) is simply a multiple of the entry cost, and cannot vary. So when profits are low, something else must change to keep the value of the firm pinned to the entry cost. This turns out to be the aggregate growth rate.

To see how this occurs, we need to consider the optimal investment behavior of the representative capitalist in response to demand fluctuations. Their investment decision determines the amount of labor allocated to research and the amount allocated to final production. To achieve this, they adjust their consumption of the final good. Their consumption is given by the residual final good production not paid to households:

$$C_t^K = Y_t - w_t L_t.$$

Combining the aggregate production function (2) and labor market clearing condition (7) yields

$$C_t^K = \tilde{\sigma}\, w_t(L_t - \sigma M_t^e\bar{\epsilon}), \tag{14}$$

where $\tilde{\sigma} \equiv 1/(1-\sigma)$. If aggregate labor supply L_t falls, the consumption of capitalists will fall without a reduction in the entry rate. To simplify the exposition, we examine the dynamics of a detrended variable, $X_t \equiv C_t^K/w_t$. This variable expresses how much of the economy's labor resource the capitalist is allocating to firm creation, and how much to production.

To see how the capitalists choose entry dynamics in equilibrium, note that with the log preferences in (4), the capitalist's Euler equation requires that their growth in consumption, denoted by $g_t^{C_K} \equiv g_t^X + g_t$, be given by

$$g_t^{C_K} = r_t^k - \rho_k$$

The expression for r_t^k is determined entirely by free entry and the firm's HJB, for a given growth rate. We show in Online Appendix B.2 that the HJB under free entry and the resource constraint in (14) characterize the dynamics of the economy in terms of labor supply L_t and two long-run stationary variables, X_t and M_t. This is summarized in the following proposition:

Proposition 1. *Entry dynamics are governed by two differential equations:*

$$X_t = \tilde{\sigma}\left[L_t - \frac{\sigma\bar{\epsilon}(\dot{M}_t+\delta_0 M_t)}{1-\delta_e},\right] \tag{15}$$

$$g_t^X = \frac{\tilde{\sigma}\bar{\omega}L_t}{\tilde{\epsilon}M_t} - \frac{(\tilde{\sigma}\bar{\omega}\bar{\epsilon}/\tilde{\epsilon}+\Phi_2+\delta_e)(g_t^m+\delta_0)}{1-\delta_e} - \delta_0 - \rho, \tag{16}$$

where g_t^X is the growth rate of X_t and g_t^m is the growth rate of the mass of firms M_t

We can see from the second equation that the key determinant of spending and growth is the ratio of labor supply to mass of firms L_t/M_t. This ratio governs the profitability of the firms. When it is high, aggregate demand for a firm's good is high relative to competition from other firms, and dividends are high. The capitalist will want to increase investment to

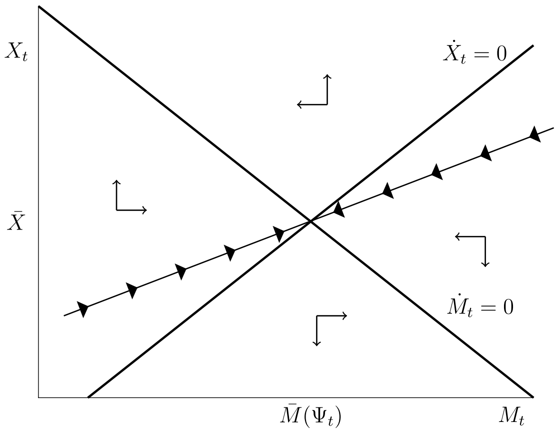

Fig. 3. Joint Dynamics of X_t and M_t.

take advantage of the high rate of profitability, and growth in productivity will rise as new firms enter. At the same time, in order to smooth consumption, the consumption of the capitalist will gradually rise faster than aggregate productivity, captured by the growth in X_t. Given the non-linearity of the equations in Proposition 1, a graphical representation is useful.

As depicted in Fig. 3, there are two stable arms of the system. The first comes from setting $\dot{M}_t = 0$ in the resource constraint in (15), and yields a linear relationship between X_t and M_t. For the mass of firms to be unchanged, the amount of entry must offset exit, requiring the capitalist to invest more of the economy's labor for entry the more firms exist at that point in time. If X_t lies above this line, the mass of firms must be falling, and the entry rate is lower than its long run value.

The second stable arm comes from setting $\dot{X}_t = 0$ in (16). On this line, the capitalist's desire to increase current consumption out of profits is exactly offset by a depreciation in the value of the portfolio from aggregate growth and creative destruction. This leads to an upward sloping linear relationship between X_t and M_t.

These stable arms meet at a unique value for $\bar{M}$ for a given amount of aggregate labor supply L, which is pinned down and increasing in the level of aggregate demand, with an associated level of detrended consumption $\bar{X}$ of the capitalist. At this point, the entry rate is set at its BGP level, just offsetting exit, and growth in productivity is equally at its long run level.

There is a unique saddle path converging to the point $(\bar{M}, \bar{X})$ in this space. Using standard arguments, we can show that the system is globally stable. In brief, for paths that begin above the saddle path, we can show that M_t reaches zero in finite time. Then consumption of the capitalist is zero, but X_t is positive, violating the feasibility constraint. For paths that begin below, the mass of firms converges to the maximum possible given the resource constraint, and consumption of the capitalist reaches zero in finite time. With log utility, this cannot be optimal.

In Online Appendix C, we present two extensions to the baseline model. The first is innovation by incumbents to improve their own product lines, as in Atkeson and Burstein (2010) and Acemoglu and Cao (2015). This fully endogenizes the innovation arrival rate ϕ. The second is allowing firms to own more than one product, as in Klette and Kortum (2004), and to add to their product portfolio through purposeful investment. This ensures that some of the creative destruction is done by existing firms. We show that under standard assumptions, in both of these extensions, *all* of the adjustment response to changes in aggregate demand are driven by entrants, supporting our empirical focus on entrant contribution to job creation in Section 2.

This section has shown that movements in aggregate demand coming from household spending changes (such as when the zero lower bound binds, inducing falls in L_t via employment rationing) can lower the returns to entrepreneurship. In turn, the entry rate falls, and growth in productivity slows. We now explore the quantitative importance of this mechanism in explaining the propagation of deleveraging shocks.

4. Calibration

In this section, we describe our calibration strategy. The unit of time is one quarter. We assume that the economy is on an initial BGP at $t = 0$, and choose parameter values so that model moments in the initial BGP match those from data. These parameter values are displayed in Table 3.

4.1. Firms

The parameters for the firm model are chosen to match key moments of firm and aggregate growth. We target an annual aggregate growth rate of 1.6% to match growth in per-capita output for the period 1980-2006. Following Foster et al. (2001) we target the BGP contribution of entry to average productivity growth to be 40%. We also target a standard deviation

Table 3
Model Parameters.

Parameter	Description	Baseline Value	Method	Source/Moment
Firm Parameters				
ρ_k	Capitalist discount rate	0.015	Calibrate	Annual R.O.E. of 7.5%
σ	Elasticity of substitution	3	Calibrate	Broda and Weinstein (2010)
ϕ	Arrival of innovations	0.92	Target	Annual stdev. incumbent growth of 1%
Δ	Innovation size	1.005	Target	Annual growth of 1.6%
$\bar{\omega}$	Avg. productivity of new products	1.30	Target	40% contr. of entry to growth
δ_0	Exogenous death rate of firms	0.010	Target	Annual entry rate of 10%
$\bar{\epsilon}$	Entry cost	187.8	Target	Avg. firm size of 19
δ_e	Innovation share on exist. product	0.59	Calibrate	Peters and Walsh (2022)
Household Parameters				
ρ_h	Household discount rate	0.01	Calibrate	
γ	Risk aversion	3.33	Calibrate	Guvenen (2006)
J	No. of labor endowment states	2	Calibrate	See Text
$(\lambda_{12}, \lambda_{21})$	Endowment switching rates	(0.5, 0.038)	Calibrate	See Text
(ℓ_1, ℓ_2)	Labor endowments	(0.418, 1.044)	Calibrate	See Text
$\underline{b}_0$	Initial borrowing limit	-1.40	Target	$r_0^h = 0$
$\underline{b}_\infty^{\text{perm}}$	Long-run borrowing limit, permanent	-1.32	Target	Output gap at $t = 8$
$\underline{b}_8^{\text{temp}}$	Minimum borrowing limit, temporary	-0.61	Target	Output gap at $t = 8$

Notes: This Table presents the calibrated parameters for the model.

of annual incumbent firm growth of 1%, matching evidence from the LBD. Together, these three moments pin down ϕ, Δ and $\bar{\omega}$ for a given value of δ_e. We set the elasticity of substitution across products to $\sigma = 3$, at the low end of the range estimated in Broda and Weinstein (2010). The discount rate of the capitalist is set at $\rho_k = 0.015$ to achieve an annual return on equity of 7.5%. The entry cost $\bar{\epsilon}$ determines the average firm size independent of the other model details, which we take to be 19 from the BDS.

The share of products created that are completely new to society, $1 - \delta_e$, we take from Peters and Walsh (2022). Together with the exogenous death rate δ_0, this determines the firm entry rate, which we calibrate to an annual value of 10% to match the BDS in the years prior to the Great Recession.

4.2. Households

Preferences. The household's discount factor is set to $\rho_h = 0.01$. We set the intertemporal elasticity of substitution (IES) for households to 1/3. This corresponds to a curvature parameter of $\gamma = 3.33$, as used in Guvenen (2006)'s two-class structure. This is also consistent with the average value reported in Barsky et al. (1997).

Labor Endowment Process. The labor endowment process is chosen to parsimoniously reflect employment risk faced by households. We set the number of labor endowments to $J = 2$. The low endowment state is 0.4 times the high-endowment state, which corresponds to the replacement ratio estimated by Shimer (2005). We normalize the scale of endowments so that $\mathrm{E}\ell = 1$, which implies aggregate labor supply is L_t. The switching intensities $\lambda_{jj'}$ are chosen so that 7% of households are in the low-endowment state and the expected duration of the low-endowment state is 6 months.

Initial Borrowing Limit. We pick the initial borrowing limit $\underline{b}_0$ so that the equilibrium interest rate is $r_0^h = 0$. Since our quantitative results are based on the model economy's response to a deleveraging shock that puts downward pressure on the interest rate, this implies that the zero lower bound binds throughout our analysis. This reflects the low-interest rate environment of the time period we are studying.[8] The calibrated value of $\underline{b}_0$ is -1.40, which corresponds to 35% of expected annual household earning in the initial BGP.

4.3. Deleveraging

We consider two demand shocks, both of which involve an unexpected tightening of the borrowing limit that occurs when the economy is in its initial BGP at $t = 0$. Similarly to Guerrieri and Lorenzoni (2017), in both cases we assume that the credit tightening occurs at a constant rate over a period of 8 quarters. Note that if the tightening is too rapid, households close to the borrowing limit would default. Dealing with endogenous default would complicate both the model and solution considerably.

In our baseline analysis we assume that the credit crunch is permanent, which leads to a persistent decrease in demand.[9] We choose the long-run borrowing limit in this case, denoted $\underline{b}_\infty^{\text{perm}}$, to match the output gap at the end of the credit

[8] Auclert and Rognlie (2018) make a similar assumption.

[9] As we show in Figure 10 of Online Appendix A.2, employment as measured by the employment to population ratio collapsed at the onset of the crisis. It showed almost no recovery while the economy remained stuck at the zero lower bound for many years, suggesting a prolonged period of weakness in

tightening period (8 quarters). We compute the output gap in the data as follows. First, we compute a baseline trend by estimating the regression

$$\log y_t = c + gt + \varepsilon_t, \tag{17}$$

where y_t is real (chained) GDP per capita, as computed by the Bureau of Economic Analysis. We use quarterly data from quarter 1 of 1947 to quarter 4 of 2019 for this regression. The output gap is the log difference between actual and predicted GDP:

$$\log y_t - (c + \hat{g}t),$$

where $\hat{g}$ is the estimated coefficient from (17). In the model, the output gap is simply the log difference between actual output and output on the initial BGP. The calibrated value of $\underline{b}_{\infty}^{\text{perm}}$ is -1.32, which represents a 5.7% tightening relative to the initial BGP.

In the second deleveraging shock we consider, the credit tightening is temporary. We assume that immediately after the tightening period is over, the borrowing limit relaxes at a constant rate back to its initial level, where it remains forever after. We choose this rate such that the borrowing limit has returned to $\underline{b}_0$ after 10 years. Given this parameterization, there is one parameter to choose: the borrowing limit at the end of the tightening period, denoted $\underline{b}_8^{\text{temp}}$. As in the permanent tightening case, we choose this parameter to match the output gap at the end of the tightening period. The calibrated value of $\underline{b}_8^{\text{temp}}$ is -0.61, which represents a 56% tightening relative to the initial BGP.

5. Results

The response of the economy to the permanent deleveraging shock is shown in Fig. 4. In the top left panel we show the calibrated borrowing limit tightening, which occurs at a constant rate over a period of 8 quarters. Aggregate labor, shown in the top right panel, drops on impact to a level below the steady state quantity implied by the new borrowing limit. Employment initially declines by nearly 10%, before gradually transitioning to its new level. Quantitatively, a shortcoming of the model is that it predicts too small and gradual of a recovery in employment after 2010.

At the same time, the entry rate decreases due to lower expected future profits. This is shown in the middle left panel. The model-implied entry rate closely mimics the sharp initial decline and subsequent gradual recovery of entry observed in the data. As shown in the middle right panel, real wages decline relative to the initial BGP after the shock due to the decrease in firm entry. The rate of decline decreases slowly, implying a significant reduction in the economy's potential on this transition. The dynamics for wage growth largely mirror the response for employment. This is consistent with our analysis in Section 3, since the equilibrium spending of households on the final good is determined entirely from rationing in the labor market. There is an initial fall in spending when demand shrinks, and households close to the borrowing constraint begin saving. The inflexibility of the real rate causes this shift to translate into a reduction in real demand, and firm profits fall. The capitalist then reduces investment in entry, and productivity and wage growth are depressed.

The nominal rigidity in wages and the zero lower bound are both crucial in delivering these results. If prices are flexible, then deleveraging has no real effects. Recall that households supply labor inelastically. With fixed nominal wage growth, there is a lower bound on the real interest rate that may prevent the bond/goods market from clearing without unemployment. If instead nominal wages are flexible, they adjust so that the real interest rate clears the bond/goods market at full employment (even in the presence of a fixed nominal interest rate). Since we do not consider exogenous productivity shocks, if employment does not change there will be no effect on real variables.

The dynamic effects of deleveraging on firm entry is further illustrated in Fig. 5. This figure shows the dynamics of the firm entry rate in response to various permanent deleveraging shocks. These shocks take the same form as the calibrated shock: the borrowing limit unexpectedly tightens for 8 quarters and then remains at that level forever after. We consider shocks in which the long-run borrowing limit is 1% and 2% tighter (labeled "-1%" and "-2%") and 1% and 2% looser (labeled "+1%" and "+2%") than the calibrated baseline, which corresponds to an 18% and 36% relative tightening. While we do not have direct evidence on borrowing limit contractions in Section 2, the interquartile range of relative house price movements for the regressions reported there is roughly 10%. Using the estimated coefficients, this would correspond to a relative fall in the establishment entry rate of 1.5% As for the overall level of the effect in our calibrated economy, the initial fall in the entry rate from 10 to 8.8% is similar to that which occurred in the aggregate data. This figure illustrates the model's ability to capture the empirical relationship between household deleveraging and firm entry shown in Fig. 2. A more severe deleveraging event is associated with a persistently lower entry rate. Note however the absence of scale effects: even though more severe tightenings lead to lower long-run aggregate labor, the entry rate always eventually recovers to that of the initial BGP.

The bottom row of Fig. 4 compares the output and consumption gaps in the model versus the data. The empirical consumption gap is computed in the same way as the output gap, using real (chained) personal consumption expenditures

aggregate spending. While our aim is indeed to explain this weakness, we do so over and above households' anticipations of persistently reduced access to credit, taking this as a fundamental shock.

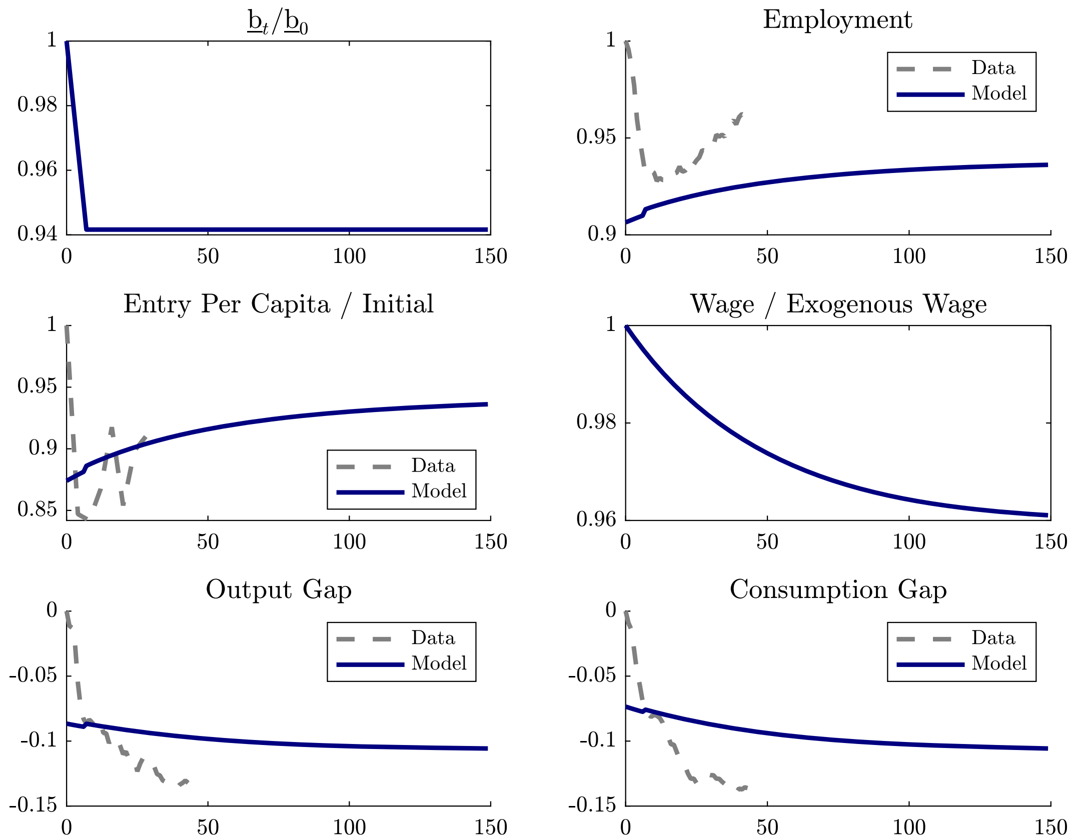

Fig. 4. Permanent Deleveraging Responses. Note: These figures show the response of the economy to the calibrated reduction in the detrended borrowing limit $\underline{b}$. Time periods are in quarters.

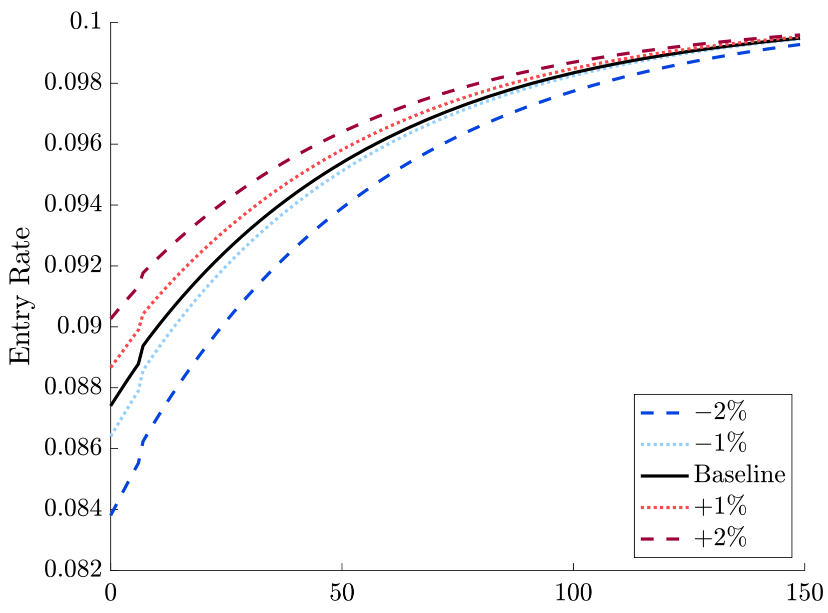

Fig. 5. Deleveraging and Firm Entry. Note: This figure shows the response of the firm entry rate for various permanent tightenings of the detrended borrowing limit $\underline{b}$. "Baseline" refers to the calibrated shock, which is chosen to match the output gap 2 years after the shock occurs. The other lines refer to permanent deleveraging shocks in which the borrowing limit is 2 percent tighter ("-2%"), 1 percent tighter ("-1%"), 1 percent looser ("+1%"), and 2 percent looser ("+2%"). Time periods are in quarters.

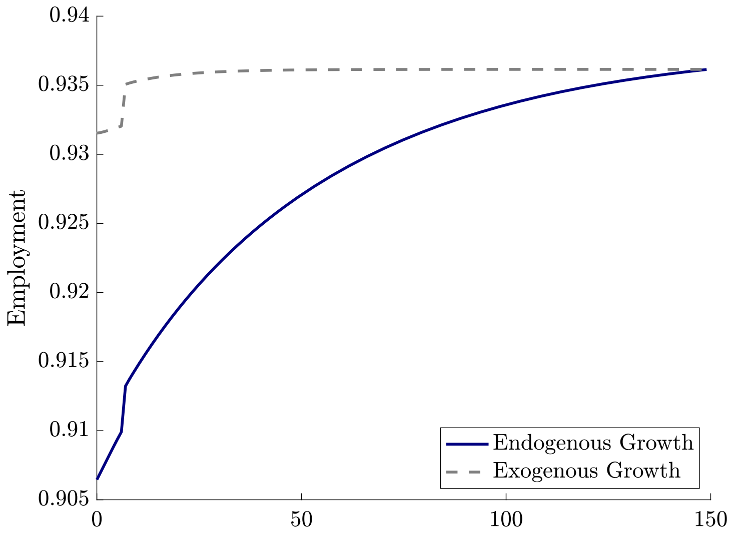

Fig. 6. Employment Dynamics With and Without Growth Feedback. Note: This figure shows the response of employment under two specifications of the baseline model. The blue line shows the response of the full model with endogenous growth. The gray line shows the response of employment if productivity growth is held fixed at its baseline BGP value and not allowed to vary. Time periods are in quarters.

per capita from the Bureau of Economic Analysis. By construction, the model output gap matches the data two years after the shock. After that time, the model somewhat underpredicts the output gap. However, similarly to the data, the output gap in this case remains persistently large. Dynamics are similar for the consumption gap.

While the reduction in demand has large consequences for the economy's potential, there is also substantial feedback from slower growth onto demand. To assess this feedback, we show in Fig. 6 how employment dynamics would evolve without endogenous growth. Namely, we consider the same tightening, but keep real wage growth fixed at its long run value of 1.6% per annum. This is represented by the gray dashed line. Without endogenous growth, the economy adjusts almost immediately to the deleveraging shock. This reflects a well-known result: by themselves, deleveraging shocks generate very little internal propagation, and the length of the aggregate response generally mirrors the length of the phase of credit tightening. Our credit tightening lasts for only eight quarters, and without endogenous growth the transitional employment dynamics are largely over by the end of this period.

With endogenous growth, the response is much longer lived. The explanation lies in the savings responses of households. When the shock hits, they find that they have under-accumulated assets, given the new predicted path for the growth rate. They know their future income will be lower than anticipated, and so want to increase savings in order to smooth consumption. Given that all agents try to save more, there is downward pressure on demand. Since the reduction in demand cannot be absorbed by falling interest rates due to the zero lower bound, employment falls further. This in turn sets off more reduction in firm profits and reduced investment in entry.

It is crucial for this mechanism to be operative that the decline in the borrowing limit be expected to be persistent. Fig. 7 shows the response of the economy to the temporary deleveraging shock. Recall that, as with the permanent deleveraging shock, this tightening was calibrated to match the output gap after 2 years. In contrast to the permanent deleveraging event (and data), in this case employment, entry, output, and consumption all recover relatively quickly to their initial BGP values.

The propagation mechanism shown in Fig. 6 is essentially absent in this case. In both the permanent and temporary deleveraging events, wage growth recovers to the initial BGP level in the long run. However, only in the permanent deleveraging case does $\overline{L}$ fall in the long run as well. As a result, there is little additional saving pressure from households who wish to smooth consumption in the face of lower future earnings. The feedback between demand and entry that was so important with a permanent shock is of only minor importance when the shock is temporary. Without this propagation mechanism, the length of the recession differs little from the duration of the credit crunch.

It is then incumbent to ask whether the credit tightening experience in the Great Recession was truly permanent, and hence whether this model provides a reasonable approximation to reality. In Fig. 8 we show two pieces of evidence that suggests credit standards have permanently tightened in the time since 2008. First we show the household debt to GDP ratio. This peaks during the Financial Crisis, and in the years since has continued to decline, despite a general recovery in the labor market and in GDP. Micro data on credit standards supports this view. In the right panel we show data on credit scores at mortgage origination over time. Credit standards tighten appreciably after 2008, and have not eased in the last decade at all. This suggests that access to credit has permanently tightened after the financial crisis, and that the permanent shock we consider in the model is the appropriate way of modeling the phenomenon.

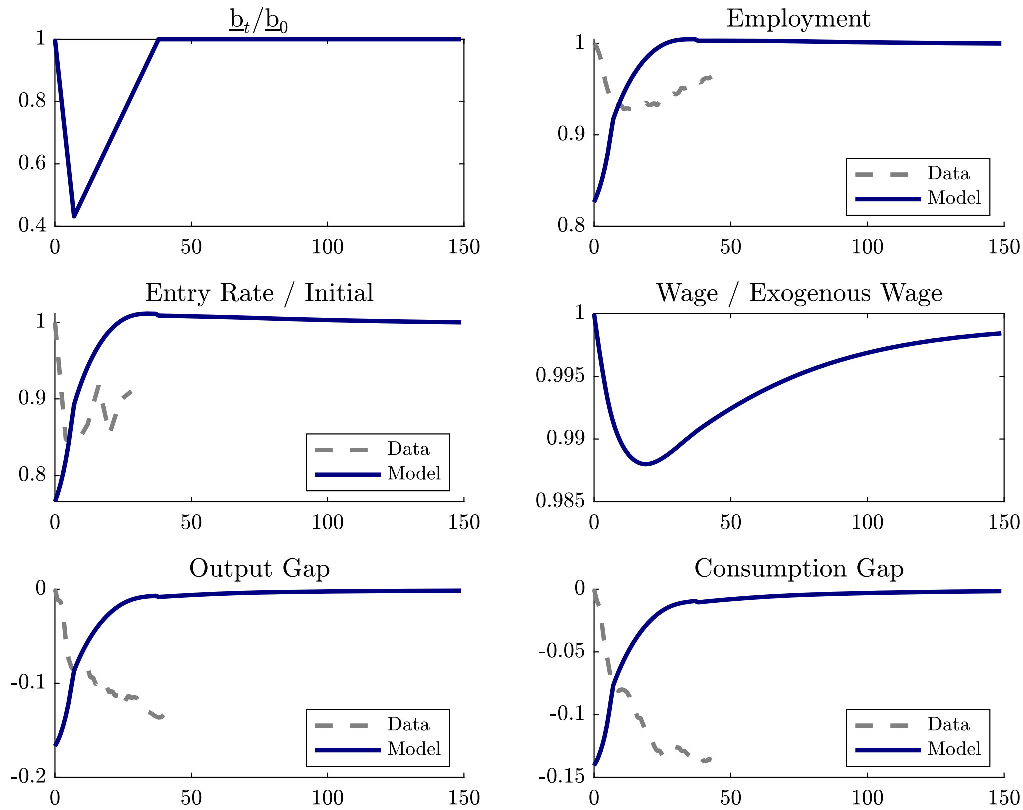

Fig. 7. Deleveraging Response: Temporary Tightening. Note: These figures show the response of the economy to the calibrated reduction in the detrended borrowing limit $\underline{b}$. Time periods are in quarters.

In Online Appendix D, we analyze an alternative calibration in which the deleveraging shock is calibrated to match data on household debt levels. This exercise provides complementary evidence that the credit tightening of the Great Recession was extremely persistent. The detrended borrowing limit remains lower than its initial level for 21 years after the initial shock, and matches the dynamics of employment and entry well. However, we impose that it does eventually revert to its pre-crisis level, and so this alternative calibration is unable to match the persistent decline in output and consumption discussed above.

In summary, the mechanism we have outlined generates a novel source of propagation of deleveraging shocks. It can help explain the slow recovery from the recent crisis, and the persistent falls in potential output. Put simply, rational pessimism about the future path for output affects spending behavior today. In a world where demand affects firm profits, this reduces investment, justifying the pessimism of households.

6. Conclusion

We have considered how aggregate demand affects productivity dynamics. Both in the data and in the theory, reductions in spending drive reduced investment in new firm and product creation. This creation is a central determinant of productivity growth. Quantitatively, a short recession implies little in the way of a productivity slowdown. However, a prolonged reduction in demand, as was seen in the decade following the 2008 crisis, can lead to a significant reduction in real wage growth. Further work could help generalize our quantitative results beyond the special deleveraging setting we have studied to more general demand and business cycle frameworks.

Data availability

Data will be made available on request.

Appendix. Supplementary material

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.red.2023.08.004.

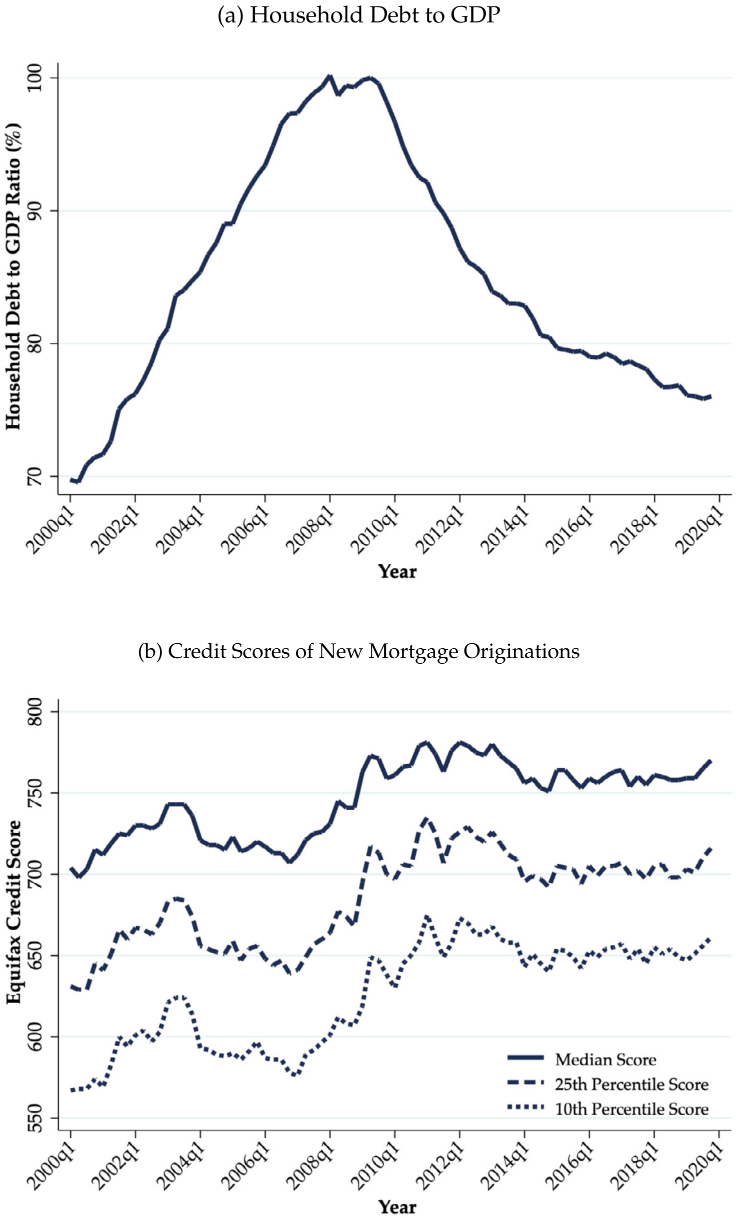

Fig. 8. Permanent Credit Tightening in the Data. Note: The top panel plots data on total Household Debt relative to GDP using data from the New York Federal Reserve. The bottom panel shows credit scores of new mortgage originations by year, using data from Equifax, accessed via the New York Federal Reserve.

References

Acemoglu, D., Cao, D., 2015. Innovation by entrants and incumbents. Journal of Economic Theory 157, 255–294.
Adelino, M., Schoar, A., Severino, F., 2015. House prices, collateral, and self-employment. Journal of Financial Economics 117, 288–306.
Aiyagari, S.R., 1994. Uninsured idiosyncratic risk and aggregate saving. The Quarterly Journal of Economics 109, 659–684.
Anzoategui, D., Comin, D., Gertler, M., Martinez, J., 2019. Endogenous technology adoption and R&D as sources of business cycle persistence. American Economic Journal: Macroeconomics 11, 67–110.
Atkeson, A., Burstein, A.T., 2010. Innovation, firm dynamics, and international trade. Journal of Political Economy 118, 433–484.
Atkeson, A., Burstein, A.T., 2019. The aggregate implications of innovation policy. Journal of Political Economy 127, 2625–2683.
Auclert, A., Rognlie, M., 2018. Inequality and aggregate demand. Tech. Rep.. National Bureau of Economic Research.
Barlevy, G., 2004. The cost of business cycles under endogenous growth. American Economic Review 94, 964–990.
Barsky, R.B., Juster, F.T., Kimball, M.S., Shapiro, M.D., 1997. Preference parameters and behavioral heterogeneity: an experimental approach in the health and retirement study. The Quarterly Journal of Economics 112, 537–579.
Benigno, G., Fornaro, L., 2018. Stagnation traps. The Review of Economic Studies 85, 1425–1470.
Berger, D., Guerrieri, V., Lorenzoni, G., Vavra, J., 2018. House prices and consumer spending. The Review of Economic Studies 85, 1502–1542.

Bianchi, F., Kung, H., Morales, G., 2019. Growth, slowdowns, and recoveries. Journal of Monetary Economics 101, 47–63.
Broda, C., Weinstein, D.E., 2010. Product creation and destruction: evidence and price implications. American Economic Review 100, 691–723.
Campbell, J.Y., Cocco, J.F., 2007. How do house prices affect consumption? Evidence from micro data. Journal of Monetary Economics 54, 591–621.
Comin, D., Gertler, M., 2006. Medium-term business cycles. American Economic Review 96, 523–551.
Davis, S.J., Haltiwanger, J.C., 2019. Dynamism diminished: the Role of Housing Markets and Credit Conditions. Tech. Rep.. National Bureau of Economic Research.
Eggertsson, G.B., Krugman, P., 2012. Debt, deleveraging, and the liquidity trap: a Fisher-Minsky-Koo approach. The Quarterly Journal of Economics 127, 1469–1513.
Fatas, A., 2000. Do business cycles cast long shadows? Short-run persistence and economic growth. Journal of Economic Growth 5, 147–162.
Foster, L., Haltiwanger, J.C., Krizan, C.J., 2001. Aggregate productivity growth: lessons from microeconomic evidence. In: New Developments in Productivity Analysis. University of Chicago Press, pp. 303–372.
Garcia-Macia, D., Hsieh, C.-T., Klenow, P.J., 2019. How Destructive is Innovation?. Econometrica 87, 1507–1541.
Guerrieri, V., Lorenzoni, G., 2017. Credit crises, precautionary savings, and the liquidity trap. The Quarterly Journal of Economics 132, 1427–1467.
Guvenen, F., 2006. Reconciling conflicting evidence on the elasticity of intertemporal substitution: a macroeconomic perspective. Journal of Monetary Economics 53, 1451–1472.
Guvenen, F., 2009. A parsimonious macroeconomic model for asset pricing. Econometrica 77, 1711–1750.
Haltiwanger, J., Jarmin, R.S., Miranda, J., 2013. Who creates jobs? Small versus large versus young. Review of Economics and Statistics 95, 347–361.
Hathaway, I., Litan, R.E., 2014. What's Driving the Decline in the Firm Formation Rate? A Partial Explanation. The Brookings Institution.
Justiniano, A., Primiceri, G.E., Tambalotti, A., 2015. Household leveraging and deleveraging. Review of Economic Dynamics 18, 3–20.
Kaplan, G., Mitman, K., Violante, G.L., 2020. Non-durable consumption and housing net worth in the great recession: evidence from easily accessible data. Journal of Public Economics 189, 104–176.
Kaplan, G., Moll, B., Violante, G.L., 2018. Monetary policy according to HANK. American Economic Review 108, 697–743.
Karahan, F., Pugsley, B., Sahin, A., 2018. Demographic Origins of the Startup Deficit. Technical Report. New York Fed. Mimeo
Klette, T.J., Kortum, S., 2004. Innovating firms and aggregate innovation. Journal of Political Economy 112, 986–1018.
Korinek, A., Simsek, A., 2016. Liquidity trap and excessive leverage. American Economic Review 106, 699–738.
Krugman, P.R., Dominquez, K.M., Rogoff, K., 1998. It's baaack: Japan's slump and the return of the liquidity trap. Brookings Papers on Economic Activity 1998, 137–205.
Lentz, R., Mortensen, D.T., 2008. An empirical model of growth through product innovation. Econometrica 76, 1317–1373.
Mian, A., Rao, K., Sufi, A., 2013. Household balance sheets, consumption, and the economic slump. The Quarterly Journal of Economics 128, 1687–1726.
Mian, A., Sufi, A., 2011. House prices, home equity–based borrowing, and the U.S. household leverage crisis. American Economic Review 101, 2132–2156.
Mian, A., Sufi, A., 2014. What explains the 2007–2009 drop in employment? Econometrica 82, 2197–2223.
Moran, P., Queralto, A., 2018. Innovation, productivity, and monetary policy. Journal of Monetary Economics 93, 24–41.
Peters, M., Walsh, C., 2022. Population Growth and Firm-Product Dynamics. Working Paper.
Pugsley, B.W., Şahin, A., 2015. Grown-up business cycles. The Review of Financial Studies.
Saiz, A., 2010. The geographic determinants of housing supply. The Quarterly Journal of Economics 125.
Shimer, R., 2005. The cyclical behavior of equilibrium unemployment and vacancies. American Economic Review 95, 25–49.
Stroebel, J., Vavra, J., 2019. House prices, local demand, and retail prices. Journal of Political Economy 127, 1391–1436.
Wolff, E.N., 2017. Household Wealth Trends in the United States, 1962 to 2016: Has Middle Class Wealth Recovered? Tech. Rep.. National Bureau of Economic Research.