Finance and Economics Discussion Series

Federal Reserve Board, Washington, D.C.
ISSN 1936-2854 (Print)
ISSN 2767-3898 (Online)

Demand Shocks and Endogenous Uncertainty

Diego Vilan

2026-050

Please cite this paper as:
Vilan, Diego (2026). "Demand Shocks and Endogenous Uncertainty," Finance and Economics Discussion Series 2026-050. Washington: Board of Governors of the Federal Reserve System, https://doi.org/10.17016/FEDS.2026.050.

Demand Shocks and Endogenous Uncertainty*

Diego Vilán†
Federal Reserve Board

Abstract

This study examines how fluctuations in firm-level uncertainty arise over the business cycle and how they influence aggregate economic activity. To do so, it develops a general equilibrium incomplete-markets model in which heterogeneous, risk-averse firms face idiosyncratic demand uncertainty and aggregate shocks to consumer credit conditions. A change in aggregate credit affects not only the expected level of firm demand, but also the cross-sectional dispersion of sales per worker by shifting the probability that firms operate at capacity. Thus, first-moment shocks give rise to endogenous second-moment effects. The model is disciplined using U.S. Compustat data on firm sales and employment and proprietary customer traffic data. The calibrated economy reproduces key cross-sectional moments and business-cycle comovements, including countercyclical dispersion in firm outcomes. Quantitatively, endogenous uncertainty accounts for roughly one quarter of the output response and one third of the employment response to aggregate credit shocks. The results suggest that uncertainty is not only an independent source of aggregate fluctuations, but also an endogenous propagation mechanism through which changes in demand conditions amplify business cycles.

Keywords: Uncertainty, Heterogeneous Firms, Cross-sectional firm dynamics.

JEL Classification: E21, E23, E32, E44.

*I owe special thanks to Vincenzo Quadrini, Juan Rubio-Ramirez, Selo Imrohoroglu, Guillaume Vandenbroucke and Pedro Silos for many insightful discussions that greatly improved this paper. I also thank participants at the Society for Computational Economics, Midwest Macro, and Dynare conferences; and seminar participants at the Federal Reserve Banks of Atlanta, St Louis and Dallas for many helpful discussions. The views expressed in this paper are solely the responsibility of the author and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System.

†Board of Governors of the Federal Reserve System, Division of Research and Statistics, 20th st and Constitution Ave. NW, Washington, DC 20551, USA. E-mail: diego.vilan@frb.gov

1 Introduction

Uncertainty fluctuations are large and strongly countercyclical (Bloom, 2014, Jurado et al., 2015). In the U.S., uncertainty has been systematically documented as having sizable adverse effects on economic activity and inflation. In terms of aggregate output, for example, Baker et al. (2024) establish that sudden changes in uncertainty may account for GDP declines in the vicinity of two percent. Gilchrist et al. (2014) report that uncertainty shocks can explain about one third of the total variation in industrial output and payroll employment; while Bachmann et al. (2013) find them responsible for manufacturing losses in excess of one percent. Moreover, Bloom (2009) and Bloom et al. (2018) argue that increased uncertainty makes it optimal for firms to wait, leading to significant declines in hiring, investment and output; and Fernández-Villaverde et al. (2015) establish that time-varying risk shocks may also have negative consequences for price stability. Related evidence from multicountry data suggests that the connection between uncertainty and real activity is not unique to the United States (Cesa-Bianchi et al., 2020).

While it has been well established that uncertainty and aggregate economic activity are negatively related, it is less evident why or how this occurs. To date most research efforts have been devoted to documenting, quantifying and understanding the effects of fluctuations in uncertainty on business conditions. In doing so, studies have often assumed the existence of sharp exogenous changes in the volatility of shocks which, mediated by physical (Bloom (2009)), financial (Gilchrist et al. (2014)) or nominal (Basu and Bundick (2017)) frictions, negatively impact mean economic outcomes. By focusing on the effects of fluctuations in uncertainty, however, almost no attention has been paid to understanding their probable sources. Recent empirical work also suggests that at least some forms of measured uncertainty behave endogenously, rising in response to adverse macroeconomic shocks rather than acting solely as independent impulses (Ludvigson et al., 2021).

Motivated by the above, this study seeks to provide evidence as to the potential origins of fluctuations in uncertainty. In doing so, it delivers a quantitative theory that is consistent with the time-varying cross-sectional properties of U.S. macroeconomic aggregates. The paper will focus on the symbiotic relationship between uncertainty and economic activity to explain how first-moment disturbances can abate or exacerbate dispersion, but also to highlight how time-varying uncertainty can affect mean equilibrium outcomes. I argue that while swings in uncertainty appear to be endogenously related to aggregate economic activity, fluctuations in idiosyncratic risk will ultimately affect macroeconomic dynamics. Intuitively recessions are times of heightened uncertainty, yet greater uncertainty may also exacerbate a recession. In particular, the study will focus on the widely held notion that consumer demand uncertainty experienced by firms could be at the heart of business cycle fluctuations. The analysis is conducted through the lens of an incomplete markets, heterogeneous-agents framework which is able to successfully reproduce the right business cycle co-movements.

The paper has two main goals. The first objective is to further our understanding of the relationship between uncertainty and mean aggregate activity. In doing so I focus on the synchronicity between uncertainty and economic outcomes, and propose an innovative channel through which the former may relate to the business environment. In particular, firms in the model face uncertainty about the number of customers they will need to serve each period (idiosyncratic), as well as the amount of resources these customers may command (aggregate). Being risk averse, firm owners will respond cautiously to changes in macroeconomic conditions, leading to cyclical employment and output fluctuations.

The second goal of this paper is to contribute to the understanding of the cross-sectional dynamics of business cycles. The availability of highly disaggregated, longitudinal microeconomic and sectoral data, has recently shed light on the idiosyncratic responses of economic agents to aggregate shocks. In turn, understanding the cross-sectional behavior of individual firms and households becomes paramount for comprehending aggregate dynamics. In the model endogenous changes in uncertainty further variations in aggregate economic activity allowing it to better replicate the observed cyclical patterns of higher moments.

Results indicate that time-varying uncertainty has significant effects on the aggregate economic activity. In the model's baseline specification, fluctuations in uncertainty account for about one-quarter of the response of output and consumption, and about one-third of the employment response, to aggregate credit shocks. Moreover, uncertainty swings act as an amplification mechanism reinforcing the original shock to mean level activity. Overall, a one percent negative shock to credit conditions leads to output and employment losses of around 0.8 and 0.6 percent respectively.

The paper makes a few additional contributions to the literature. First, it introduces an innovative way of modeling fluctuations in consumers' demand. Rather than assuming exogenous changes to a household's discount factor, the model will keep track of the distribution of customers visiting a firm. Second, the proposed framework sheds light on the relationship between uncertainty and risk-averse behavior, in that higher perceived risk might exacerbate the effects of first-moment disturbances hitting the economy. Lastly, the study proposes a parsimonious framework capable of capturing fluctuations in uncertainty which requires no nominal rigidities and offers a tractable closed form solution.

1.1 Related Literature

This study is closely related to a fast-growing body of literature studying the effects of time-varying uncertainty on economic activity. It follows Bloom (2009), Basu and Bundick (2017), and Leduc and Liu (2016) in that fluctuations in second moments have first order aggregate effects. The overriding idea in this area of research is that spikes in uncertainty, channeled through some adjustment friction, generate the observed fluctuations in economic activity. For broad surveys and measurement contributions, see Bloom (2014), Jurado et al. (2015), Baker et al. (2016), and Scotti (2016). Moreover, the paper also relates to the scholarly research focusing on uncertainty fluctuations as an endogenous outcome rather than a cause. Recent empirical work by Ludvigson et al. (2021) sharpens the distinction between uncertainty as an exogenous impulse and as an endogenous response to business-cycle shocks, while Orlik and Veldkamp (2024) emphasize the role of black-swan states in shaping uncertainty shocks. In this view, Bachmann and Moscarini (2011) propose a model in which recessions tend to incentivize firms' risk taking behavior and hence lead to higher cross-sectional dispersion. Similarly, Fostel and Geanakoplos (2012) and D'Erasmo and Boedo (2011) suggest alternative mechanisms capable of generating countercyclical uncertainty.

The proposed framework also represents a natural extension to Bewley-type models such as Aiyagari (1994), Huggett (1997) and Krusell and Smith (1998). These models introduce idiosyncratic risk into an incomplete markets neoclassical framework, but focus on labor-income risk, rather than demand uncertainty. Furthermore, the paper closely follows Angeletos (2007) and Quadrini (2017), both of which provide the theoretical underpinnings behind the set-up as well as the chosen solution method.

The study also relates to the literature seeking to understand the idiosyncratic effects of aggregate shocks. Higson et al. (2002) and Higson et al. (2004) report that rapidly growing and rapidly declining firms appear to be less sensitive to negative macroeconomic disturbances relative to those firms in the middle range of growth. This appears to be consistent with the fact that the higher moments of the distribution of firm growth rates have significant cyclical patterns. Similarly, Kehrig (2011) finds that the cross-sectional dispersion of firm-level total factor productivity in the U.S. tends to be greater in recessions than in expansions. More recently, Dew-Becker and Giglio (2023) construct a market-based measure of cross-sectional uncertainty from option prices and document its own cyclical dynamics, while Herskovic et al. (2020) link firm-level volatility to granular network exposures.

In terms of production, some papers assign a productive role to consumer demand for goods and services. With this in mind, this study follows Bai et al. (2025) and Petrosky-Nadeau and Wasmer (2015) in that output will not only be a function of factor inputs (like in any neoclassical framework), but consumer demand will play a paramount role in determining the level of economic activity. Moreover, in line with Arellano et al. (2019) the framework also explores the effects of input pre-commitments in optimal firm behavior.

Finally, the study is also related to the literature highlighting the effects of financial frictions on the interaction between uncertainty and economic activity. Gilchrist et al. (2014) argue that increases in firm risk lead to higher bond premia and a higher cost of capital, which in turn, triggers the prolonged decline in investment activity. It also follows Jermann and Quadrini (2012) in that the financial sector may be the source of the business cycle and not solely a propagation channel for shocks that hit other sectors of the economy.

The remainder of this paper is organized as follows: Section 2 presents the empirical motivation and analysis from the Compustat and consumer traffic data sets. Section 3 explains the model, and 4 describes its calibration. Finally, Section 5 presents the main results while

Section 6 draws some final conclusions.

2 Empirical Motivation

2.1 Time series facts

The negative association between uncertainty and economic activity finds substantial empirical support in the U.S. economy. The above patterns, however, are not exclusive to it and a plethora of studies have recorded similar realities in countries around the globe. Bachmann et al. (2013) use German data to provide evidence as to the detrimental effects of uncertainty in that country. For the UK, Denis and Kannan (2013) estimate that uncertainty shocks generate industrial production and output losses, while Bloom et al. (2007) find evidence that supports the claim that higher uncertainty reduces domestic firms' capital expenditures. Similar conclusions have been reached for developing economies. Arslan et al. (2015) establish that a one standard deviation increase in aggregate uncertainty generates a 4 percent drop in Turkey's GDP growth rate; while Fernández-Villaverde et al. (2011) compute the negative effects of interest rate volatility for a group of Latin American economies. Globally, Baker et al. (2016) document the effects of uncertainty in slowing down the global recovery.

Given its intrinsically unobservable and yet broad nature, uncertainty can be very hard to measure. It reflects the ambivalence in the minds of consumers, investors, and policymakers about the likelihood of potential future outcomes. It can also reflect skepticism about aggregate events such as the growth rate, credit conditions and exchange rates; or micro phenomena such as industry level legislation or personal ambiguity. Not surprisingly, a plethora of proxies have been developed over the last years in an attempt to capture sudden variations in risk. One of these measures is the CBOE Volatility Index (VIX) which captures the expected thirty-day-ahead implied volatility backed out from option prices. An alternative proxy for uncertainty is the corporate bond spread computed as the difference between the Baa 30-year yield and the U.S. Treasury yield at a comparable maturity. Another measure frequently used is the disagreement amongst professional forecasters. Periods of higher uncertainty usually correlate with greater dispersion in professionals' opinions. The intuition is that uncertainty makes it harder for agents to make accurate predictions. Finally, Baker et al. (2016) develop an alternative proxy for uncertainty by recording the frequency of newspaper articles reporting on this topic. Figure 1 plots a selection of commonly used empirical measures of uncertainty over the business cycle.

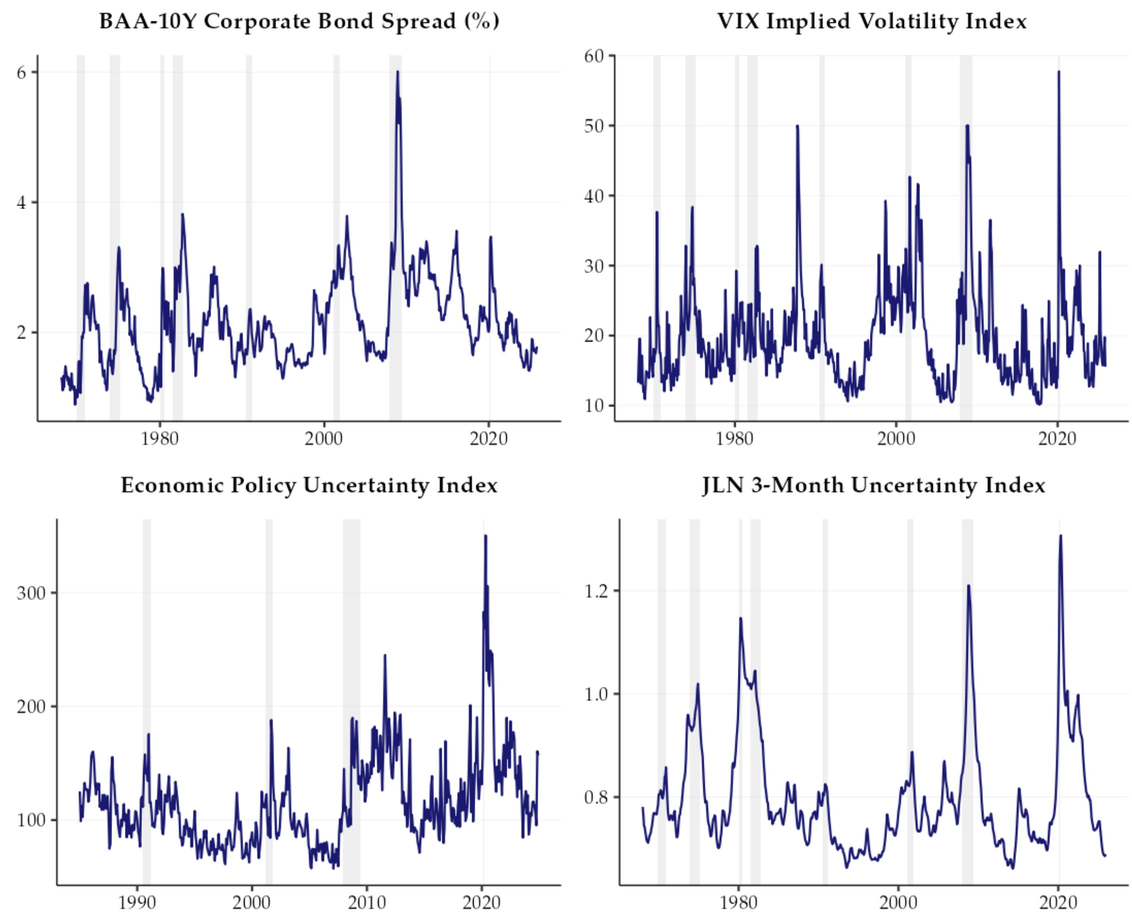

Figure 1: Uncertainty indicators over the Business cycle.

Regardless of which metric is used, virtually every indicator of uncertainty rises in recessions and subsides during expansions. Conversely, measures of economic activity tend to move in communion with the cycle. Figure 2 shows this graphically, plotting the business cycle evolution of six macroeconomic indicators. Intuitively as economic activity slows down, jobs are lost, consumption falls and capacity utilization rates plummet. Additionally, as aggregate credit conditions deteriorate, sales growth slows down and corporate net worth suffers. This is the negative association between uncertainty and economic activity which will be at the core of this study. In particular, by focusing on consumer demand uncertainty the model will successfully reproduce the business cycle dynamics in all six macroeconomic yardsticks mentioned above.

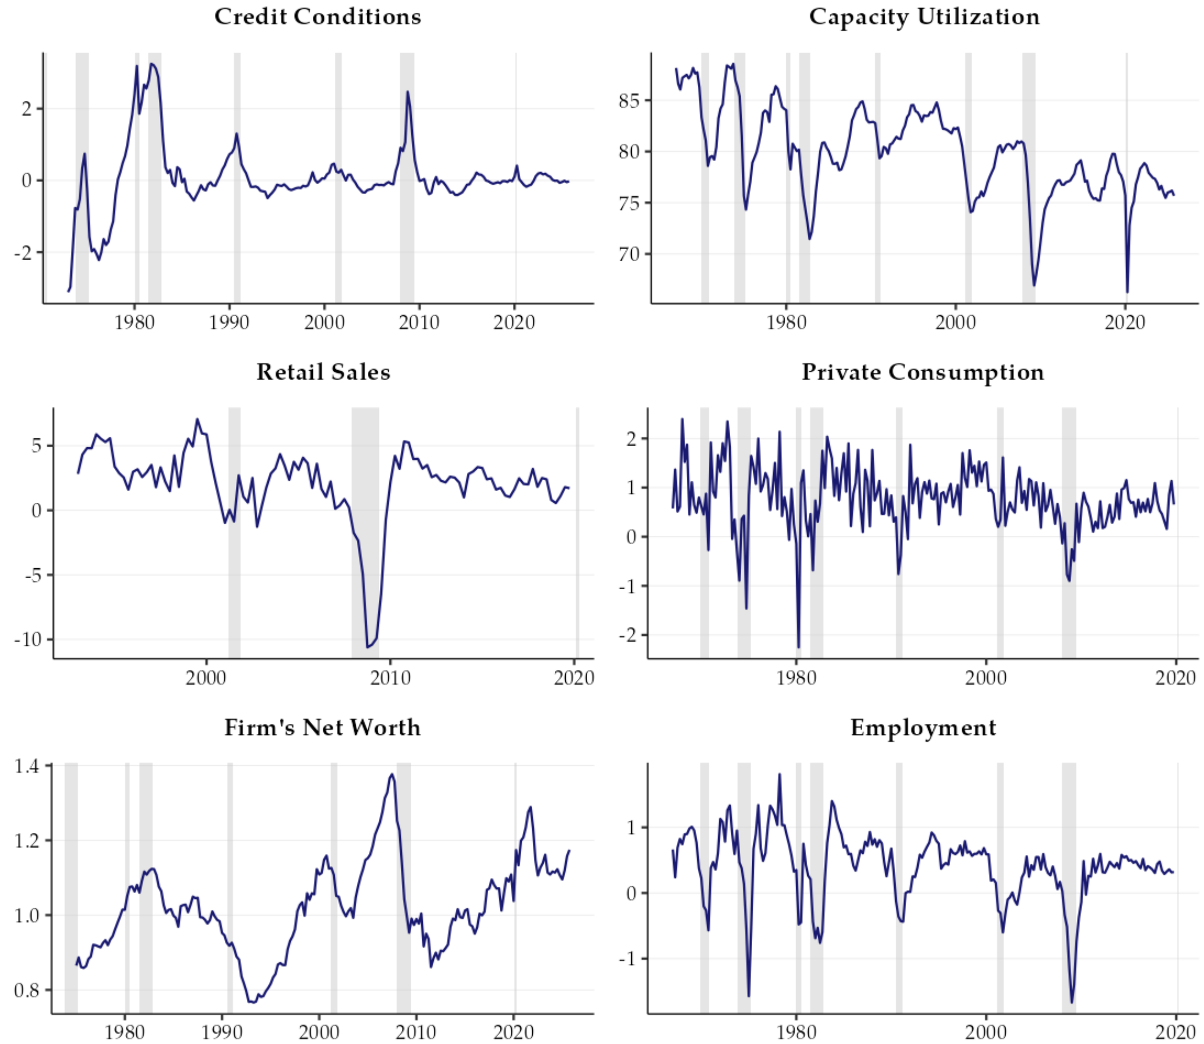

Figure 2: Uncertainty and Economic Activity. *Consumption* corresponds to the year-over-year changes in Personal Consumption Expenditures (PCE) as recorded by the BEA, while *Employment* tracks the year over year changes to the level of total non-farm, quarterly employment. *Capacity Utilization* refers to the percentage of industrial capacity currently being used by firms domestically to produce the demanded finished products as compiled by the Board of Governors of the Federal Reserve System. *Retail Sales* correspond to the yearly change in the level of retail and food services sales as measured by the U.S. Census Bureau, and Credit Conditions refer to the Federal Reserve Bank of Chicago's National Financial Conditions Index (NFCI), where positive values of the index indicate that financial conditions are tighter than average. Finally, *Firms' Net Worth* tracks the evolution of the non-financial corporate business sector's net worth as a percentage of GDP.

2.2 Firm-level facts

Researchers focusing on the impact of uncertainty on individual firms and households have found that uncertainty at the firm level is also negatively associated with growth and economic activity. Kehrig (2011), for example, shows that uncertainty about plant-level TFP among U.S. durable goods manufacturers rises sharply during recessions, affecting firms' entry and survival rates. Vavra (2014) establishes that price uncertainty also surges in recessions, creating additional challenges for the conduct of monetary policy. Similarly, Higson et al. (2002) find that firm-level risk is negatively correlated with the business cycle and affects firms unevenly, while Leahy and Whited (1996) document a strong negative relationship between uncertainty and investment among publicly listed U.S. firms.

More recent work reinforces this evidence by using direct and more granular measures of firm-level uncertainty. Altig et al. (2022) construct measures of subjective business uncertainty from survey-based probability distributions over firms' own sales, employment, and investment outcomes. Using related Census microdata, Bloom et al. (2022) find that higher subjective uncertainty is associated with lower investment, employment growth, and shipments growth. In the same spirit, Alfaro et al. (2024) show that uncertainty shocks reduce firm investment, with stronger effects among financially constrained firms, while Hassan et al. (2019) document that firms exposed to higher political risk reduce both hiring and investment. Taken together, this evidence suggests that uncertainty is not only countercyclical in aggregate data, but also shapes real decisions at the firm level.

The primary firm-level data source used in this paper is the US Compustat database. Compustat North America provides the annual and quarterly Income Statement, Balance Sheet, Statement of Cash Flows, and supplemental data items on most publicly held companies in the United States and Canada. Financial data items are collected from a wide variety of sources including news wire services, news releases, shareholder reports, direct company contacts, and quarterly and annual documents filed with the Securities and Exchange Commission. Compustat files also contain information on aggregates, industry segments, banks, market prices, dividends, and earnings. Depending upon the data set, coverage may extend as far back as 1950 through the most recent year-end.

Using Compustat has some advantages versus using census data sets like the Longitudinal Research Dataset (LRD) or the Annual Survey of Manufacturers (ASM), because firm-level data are accessible to all researchers in different countries, and the panel for the US goes as far as the 1950s. Naturally, these data are not without flaws, the most commonly recognized being the fact that the firms recorded in Compustat account for about one-third of US employment (Davis et al. (2006)).

The data set comprises of over forty years of data (1980-2024), with cross-sections that have, on average over 7,000 firms per year. From the original Compustat data, I select firms that report information on gross and net sales, employment and capital stocks. Following Bloom (2009) I drop firms with missing information as well as remove outliers. To calculate firm-level employment growth rates I use the symmetric adjustment rate definition proposed in Davis et al. (2006):

$$g_{h,t} = \frac{h_t^i - h_{t-1}^i}{0.5 * (h_t^i + h_{t-1}^i)}$$

Firm-level sales growth rates are simple log-differences. To focus on idiosyncratic changes that do not capture differences in industry-specific responses to aggregate shocks, I follow Bachmann et al. (2013) in removing firm effects from employment and sales growth rates. Annual GDP and inflation data come from the Federal Reserve Economic Data (FRED) database. All moments are robust to different inflation-index specifications. Table 1 summarizes some of the statistical properties of the US Compustat data set.

Table 1: U.S. Compustat Moments 1980-2024

	ln(sales)	ln(emp)
Cross-sectional dispersion	2.619	2.373
Cross-sectional skewness	-0.319	-0.053
Cross-sectional kurtosis	2.894	2.386
Dispersion growth rate corr w/ cycle	-0.194	0.156
Skewness corr w/cycle	-0.007	0.006
Kurtosis corr w/cycle	-0.206	-0.007
# observations (ave. per year)	8,002	7,027
# observations (total)	360,107	316,227

Source: Compustat and U.S. Bureau of Economic Analysis

Table 1 documents the presence of significant deviations from symmetry and normality in the data, with average log-sales dispersion of 2.619 and average log-employment dispersion of 2.373. At the same time, the log distributions are relatively well behaved. Log sales exhibit only mild negative skewness, while log employment is close to symmetric; kurtosis is close to, or below, the normal benchmark of 3, indicating that the log-transformed and trimmed distributions do not display pronounced fat tails. Thus, although sales and employment in levels are strictly positive, highly dispersed, and strongly right-skewed, taking logs substantially regularizes the cross-sectional distributions.

These empirical features motivate the lognormal specification used in the calibration. The assumption is not that the full empirical distribution is exactly lognormal, especially in the extreme tails, but rather that lognormality provides a parsimonious way to represent the observed cross-sectional dispersion of positive and highly heterogeneous variables. Since the model is primarily disciplined by average cross-sectional dispersion, the lognormal specification offers a consistent reduced-form approximation for firm sales and employment, in the same spirit as the customer-traffic calibration.

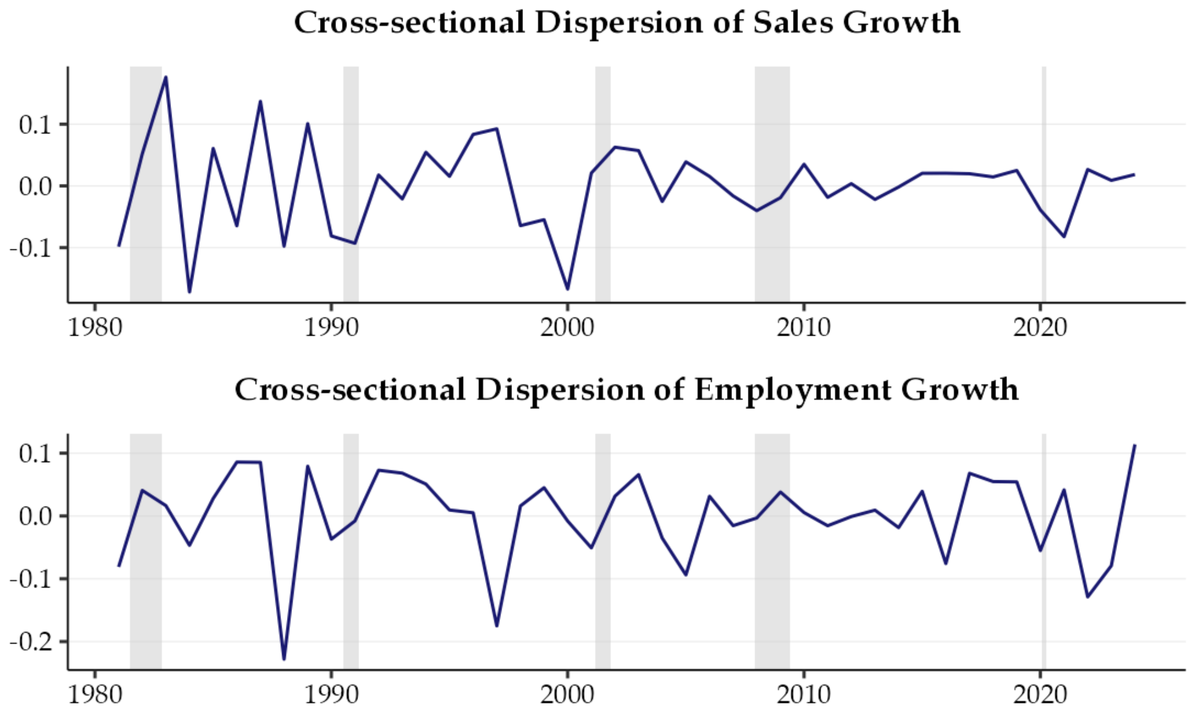

Figure 3: Uncertainty over the business cycle

The data also point to substantial cross-sectional heterogeneity in firm-level growth rates, and this heterogeneity varies systematically over the business cycle. In particular, the cross-sectional dispersion of sales growth is negatively correlated with aggregate economic activity, indicating that sales-growth dispersion tends to rise during downturns and fall during expansions. The corresponding pattern for employment-growth dispersion is weaker and mildly procyclical in this sample. Figure 3 plots the evolution of the cross-sectional dispersion of sales and employment growth rates over the course of five NBER recessions, illustrating the extent to which firm-level outcomes become more dispersed around periods of aggregate stress. These patterns suggest that aggregate downturns are not associated only with lower average growth, but also with a broader distribution of firm-level outcomes. In recessions, some firms experience sharp declines in sales and employment, while others are less affected or even continue to expand. This widening of the cross-sectional distribution is a central feature of the data that the model seeks to capture.

These observations are consistent with a broader empirical literature documenting countercyclical dispersion in firm-level outcomes. Bloom (2009) constructs measures of microeconomic uncertainty using firm-level dispersion in Compustat profit growth and stock returns, while Bachmann and Bayer (2014) document related cyclical patterns in the dispersion of firm-level outcomes using German firm data. Similarly, Kehrig (2011) finds countercyclical dispersion in plant-level productivity. More directly related to the sales-growth moments reported here, Higson et al. (2002) show that the cross-sectional variance of sales growth among U.S. quoted firms is negatively correlated with aggregate activity, and Bloom et al. (2018) document that firm-level sales-growth dispersion rises in recessions. These stylized facts provide empirical regularities to be matched by the proposed framework.

3 The Model

The baseline model builds on elements of Quadrini (2017) and features two sectors: an entrepreneurial sector and a household sector. Entrepreneurs are the sole owners of firms and are responsible for producing goods in the economy. Households supply labor and demand consumption goods. Firms face uncertain demand for their products and hold financial assets to mitigate the effects of adverse idiosyncratic shocks. The full set-up is described below.

3.1 Households

Time is discrete, indexed by $t \in \{0, 1, ..., \infty\}$. There is a continuum of infinitely-lived households whose preferences are separable in consumption, c_t, and labor supply, h_t, as described by:

$$U_H = E_t \sum_{t=0}^{\infty} \beta^t \left(c_t - \gamma \frac{h_t^{1+\tau}}{1+\tau} \right)$$

where E_0 is the conditional expectation operator, β is the discount factor, $\gamma > 0$ measures the relative disutility of labor effort and $\tau > 0$ is related to the Frisch elasticity of labor supply. Households supply labor in a competitive market and allocate their labor and financial earnings between consumption goods and risk-free assets. Their budget constraint is:

$$w_t h_t + \frac{b_{t+1}}{R_t} \geq c_t + b_t$$

where $w_t h_t$ is the period real labor income, R_t is the gross interest rate and b_{t+1} is the loan contracted in period t and due in period $t+1$. Balances are settled every period and there is no default. Households may accumulate intertemporal assets, but face the following borrowing constraint:

$$\Omega \geq \frac{b_{t+1}}{R_t} \quad \forall t$$

Households will seek to purchase consumption goods before collecting their labor income. Since the goods are acquired before wages are paid and before the opening of financial markets for inter-temporal transactions, all purchases are paid with intra-period credit. This intra-period credit is subject to a limit θ_t, which is stochastic and follows the process:

$$\begin{aligned} \ln \theta_t &= \rho \ln \theta_{t-1} + \varepsilon_t \\ &: \quad \varepsilon_t \sim N(\mu_\varepsilon, \sigma_\varepsilon^2) \end{aligned}$$

This time-varying limit is meant to capture the evolution of aggregate consumer credit conditions in the economy.

3.2 Entrepreneurs

There is a continuum of entrepreneurs indexed by i with lifetime preferences over consumption streams given by:

$$U_E^i = E_0 \sum_{t=0}^{\infty} \beta^t \ln c_t^i$$

where E_0 is the expectation operator conditional on the information available at $t = 0$ and β the discount factor. Entrepreneurs are individual owners of firms and produce a homogeneous, non-storable and competitively traded consumption good. Firms have revenue functions $y_t^i h_t^i$, where the variable h_t^i is the labor input and y_t^i represents output per worker at time t. A firm's level of output per unit of labor is an idiosyncratic stochastic variable that will be defined below. For the moment what matters is that the operation of every firm is subject to an idiosyncratic shock y_t^i.

Following Arellano et al. (2019) it is assumed that entrepreneurs choose the input of labor before observing the actual realization of y_t^i. Moreover, I assume that the wage rate cannot be made contingent on the realization of the idiosyncratic uncertainty. Since labor markets are competitive, this implies that wage rate will be the same for all firms. Markets are assumed to be incomplete, with only one asset available for entrepreneurs to self-insure against the idiosyncratic risk: a non-contingent bond b_t^i that pays the gross interest rate R_t. The entrepreneur's budget constraint is therefore:

$$y_t^i h_t^i + b_t^i \geq c_t^i + w_t h_t^i + \frac{b_{t+1}^i}{R_t}$$

where h_t^i is the labor input provided by households to firm i in period t, and y_t^i is firm $i's$ idiosyncratic output per worker. All in all, these assumptions imply that the firm faces a risk in the choice of labor which cannot be fully insured.

3.3 Production and Demand Uncertainty

Every period consumers get randomly distributed among producers. In particular, assume that each household visits $\chi < 1$ producers. Even if each household visits the same number of producers, the distribution of consumers over producers is not uniform. This implies that some producers will receive more consumers (per-unit of labor) than others. As such, the demand uncertainty faced by firms derives from the randomness in which households get distributed among entrepreneurs.

Denote by n_t^i the number of consumers per unit of labor received by producer i in period t. This variable is stochastic with probability density $f(n)$. Since each household visits χ producers, the distribution must satisfy $\int nf(n)\,dn = \chi$. That is to say, the average number of consumers per worker received by each producer is χ.

Given the choice of h_t^i, a firm can produce at most $\bar{y}h_t^i$, where $\bar{y} > 0$ is a constant and represents a technological constraint. Since all entrepreneurs utilize the same production technology, $\bar{y}$ will be the same for all firms. The quantity $\bar{y}h_t^i$ represents the firm's period t production capacity after hiring h_t^i units of labor. The actual production, however, depends on the quantity of goods that the firm can sell, which is unknown to the entrepreneur at the time he or she must make the hiring decision.

Each period can be thought of being divided into three subperiods. In the first subperiod firms choose employment h_t^i and promise to pay workers the wage w_t. In the second subperiod households visit producers to shop for consumption goods and engage in production. In the third subperiod households are allowed to re-trade the goods acquired from the entrepreneurs in a Walrasian market and all credit/debit positions, including the promised wages, are settled. Each subperiod is outlined below.

Subperiod 1: Hiring stage. Entrepreneurs hire labor h_t^i and set their period productive capacity $\bar{y}h_t^i$. The hiring decision takes into account the uncertainty about the goods that the firm will actually be able to sell in the second subperiod.

Subperiod 2: Decentralized shopping and production. Since a household has a credit capacity of θ_t and visits χ firms, the spending capacity in each producer is θ_t/χ. Therefore a firm that receives $n_t^i h_t^i$ consumers can sell at most $n_t^i h_t^i \theta_t/\chi$ units of goods, that is the number of consumers multiplied by the credit capacity of each consumer. Assuming that producers have all the bargaining power, the revenue per worker of firm i is:

$$y_t^i = \begin{cases} \bar{y} & \text{if } \; n_t^i \left(\frac{\theta_t}{\chi}\right) \geq \bar{y} \\ \\ n_t^i \left(\frac{\theta_t}{\chi}\right) & \text{if } \; n_t^i \left(\frac{\theta_t}{\chi}\right) < \bar{y} \end{cases}$$

Hence production per unit of labor will be determined by the number of customers that a firm receives, as well as by their purchasing capacity (intra-period credit). Last, sales for a firm that hires h_t^i workers is:

$$Y_t^i \;\; = \;\; y_t^i h_t^i$$

The assumption that the producers hold all the bargaining power guarantees that, when the demand is smaller than the production capacity of the firm, the firm does not sell to customers more goods than their credit capacity. At the same time, the assumption that households are allowed to re-trade the acquired goods in subperiod 3 (as described below) guarantees that the firm does not charge an interest rate on the intra-period credit when the demand exceeds the production capacity of the firm. Notice that charging an interest rate is equivalent to charging a higher price for the good (units of consumption goods in subperiod 3 per one unit of consumption goods in subperiod 2).

Subperiod 3: Centralized trading and settlements. Since during the second subperiod households are randomly matched with producers, the quantity of goods purchased differs across households. By assuming that at this stage the acquired goods can be re-traded in a centralized, anonymous market, all households face the same optimization problem at the end of the period. Specifically, they solve the recursive problem below:

Let $S_t = \{B_t, \theta_t\}$ represent the aggregate states of the economy at time t, namely the extent of credit conditions in the economy and the aggregate level of wealth[1]. Recursively, the household's optimization problem can be stated as:

$$
\begin{aligned}
V(S,b) &= \max_{c,h,b'} \left\{ c - \gamma \frac{h^{1+\tau}}{1+\tau} + \beta E_\theta V(S',b') \right\} \\
s.t. &: \ wh + \frac{b'}{R} \geq c + b \\
&: \ \Omega \geq \frac{b'}{R}
\end{aligned}
$$

Their optimal policies satisfy the first order conditions:

$$
\begin{aligned}
\gamma h_t^\tau &= w_t \\
u_c(c_t, h_t) &\geq \beta R_t E_t u_c(c_{t+1}, h_{t+1})
\end{aligned}
$$

where the last condition will be satisfied with equality if the inter-temporal borrowing constraint is binding.

Overall, the model's timing is as follows: each entrepreneur i enters period t with risk-free bonds b_t^i and chooses the labor input h_t^i knowing θ_t but before the realization of the idiosyncratic matching n_t^i takes place. Labor markets are competitive and the real wage w_t fluctuates to equate demand and supply. Once n_t^i is known production takes place, consumers acquire goods on credit, and firms' profits are realized. Following households collect their wages and balances are settled. In settling their liabilities, households may choose to re-trade some of their purchased goods in an anonymous Walrasian market which opens at the end of every period. Agents who acquired goods on credit beyond their actual possibilities might seek to sell some of their purchases to settle claims. Similarly households who were not able to purchase enough goods from the firms they were matched with, might seek to increase their consumption via this market. Finally, each agent chooses the next period's bond holding b_{t+1}^i.

[1] I define B_t as $B_t = B_t^E + B_t^H$ where $B_t^E = \int b_t^i \ dF(i)$ and $B_t^H = b_t$ in equilibrium.

Figure 4 schematically represents the model's timing.

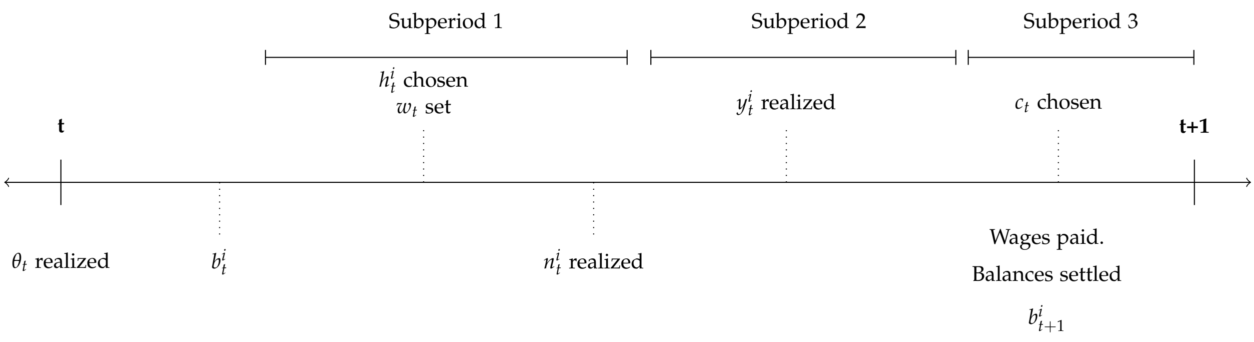

Figure 4: Model's Timing

3.4 Endogenous Uncertainty

Every period entrepreneurs must decide on their optimal level of output and employment. They must do so aware of the state of aggregate credit conditions in the economy, but before knowing the actual number of customers that will visit their store. To make their decision, entrepreneurs will take into account their current level of assets b_t^i and form expectations about their future level of sales. Firms will base these forecasts on the probability distribution of n_t^i conditional on the realization of θ_t. This conditioning is relevant since the level of aggregate credit will have first and second moment effects on the distribution of sales per worker as detailed below.

Since the realization of θ_t represents an aggregate shock and, given the definition of sales in the model, positive realizations of this variable will shift the distribution of sales per unit of labor to the right, while negative ones will do so to the left. This represents a first moment effect on the distribution of sales, implying a higher or lower mean, yet a constant level of dispersion. The intuition is simple, when aggregate credit conditions in the economy are good, agents are able to demand more goods and firms expect their average period sales to be higher. The opposite happens during a contraction. The figure below describes the effects of a positive increase in the level of aggregate credit on the distribution of sales per worker.

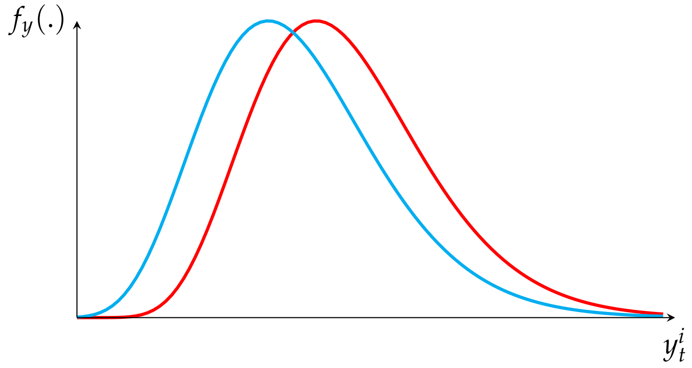

Figure 5: Distribution of sales per worker

The realization of θ_t, moreover, will also influence the level of idiosyncratic uncertainty faced by individual producers in the economy. Given that workers are subject to a technological constraint $\bar{y}$, the aggregate level of credit in the economy will condition the maximum number of clients that each worker can care for. Assume for example that capacity is set to $\bar{y} = \Omega$ units of the consumption good. A worker in this economy could sell all Ω units to a sole customer with enough credit limit to demand the worker's entire output, or Ω/n units to n customers with lower credit limits. Effectively this constraint will produce a censored distribution of clients per worker as described in Figure 6:

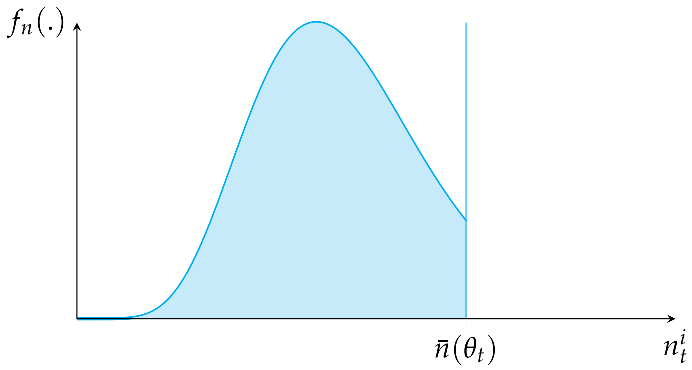

Figure 6: Distribution of customers per worker

A low realization of θ_t means that each worker can serve a large number of clients, yet also implies it will require a substantial number of clients for that worker to be profitable. This increases the level of risk per worker borne by the entrepreneur. Conversely, a high realization of θ_t suggests that workers will need to serve fewer customers to become profitable, since each one will demand a greater number of goods, effectively lowering the risk per worker. Figure 7 highlights the change in dispersion for an increase in the aggregate level of credit from θ_t^1 to θ_t^2, where $\theta_t^2 > \theta_t^1$. The shaded area represents the reallocation of probability mass into the new censoring point, and hence the overall reduction in the level of uncertainty per unit of labor that any firm must sustain when hiring a worker.

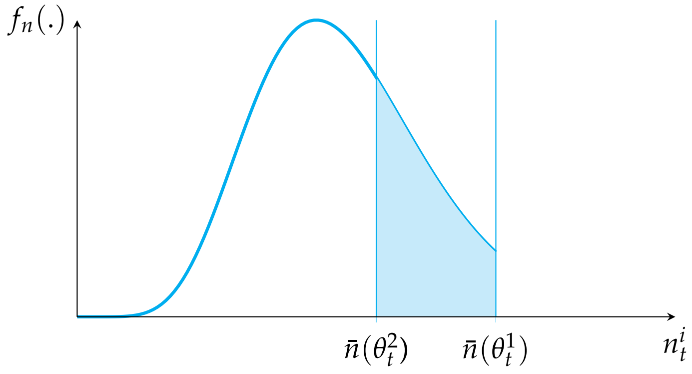

Figure 7: $\theta_t^2 > \theta_t^1$

Overall these two features capture the effects of the aggregate state of the economy in shaping entrepreneurs' expectations about their future level of sales and in doing so condition their production and hiring decisions. As credit conditions improve, the average expected level of sales rises (first moment effect) enticing firms to hire more workers. At the same time, the variance of the distribution over which agents form their expectations decreases (second moment effect), further increasing entrepreneurs' demand for labor. Figure 8 represents the combined first and second moment effects, where the shaded area represents the decrease in dispersion induced by an improvement in aggregate credit conditions.

In an expansion, higher expected sales coupled with a lower risk per unit of labor leads entrepreneurs to revise their production plans and increase their hiring. The opposite hap-

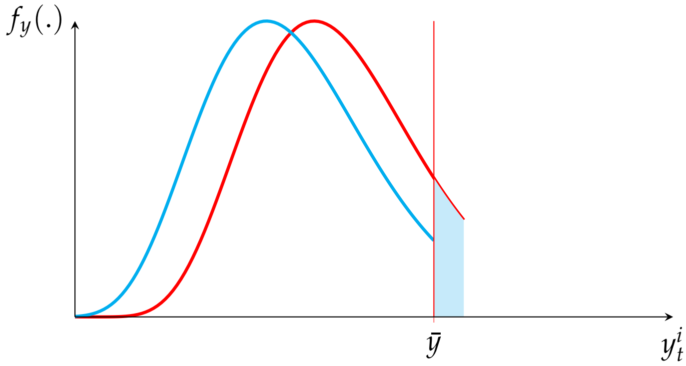

Figure 8: Endogenous Uncertainty

pens in a contraction. The final level of production, nonetheless, will also depend on an entrepreneur's ability to hedge the total production risk as described in the next section.

3.5 Risk and Return trade-off

The level of aggregate credit in the economy will condition both, the expected average level of sales and the level of risk per unit of labor. It will also limit the number of customers that may be cared for per worker. Beyond $\bar{n}$ an entrepreneur knows that those clients will not be served and revenue will be lost. Hiring more employees allows business owners to serve more customers, but also increases the overall level of risk they must bear. Given that workers collect their wages independently of the achieved level of sales, the bigger the wage bill the greater the entrepreneur's exposure to an adverse realization of n_t^i. In turn, as more workers are hired the firm's expected sales rise, but so does the size of a potential loss. This represents the fundamental trade-off solved by entrepreneurs when confronted with the task of choosing their optimal level of inputs.

This trade-off has two principal components: the distribution of sales per firm as well as the size of its wage bill. In terms of sales, the relevant underlying distribution is that of customers per firm. As more workers are recruited, this distribution achieves a higher mean and a higher variance. This implies that the firm's expected sales will increase, but so will

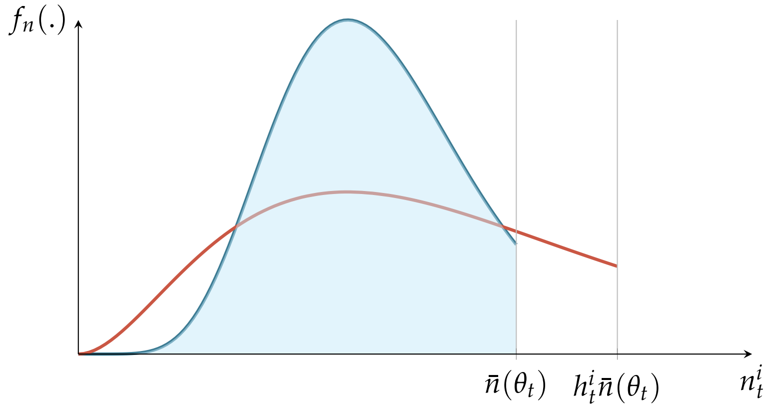

Figure 9: Distribution of customers per firm

the level of risk faced by its owner. Figure 9 sketches the mentioned changes endured by the distribution of customers for a firm which increases its number of employees.

The second element affecting the entrepreneur's profitability, the wage bill, also increases as more workers are recruited. Crucially, while the wage bill increases monotonically with every new employee, the probability of additional customers does not. Figure 10 simulates the profit distribution for three different firm sizes: 2, 4 and 6 employees. As the number of employees grows the resulting distribution has a higher mean and higher variance, yet more importantly, it begins to increasingly gain mass on the low outcome events. As such, even when the shape of the customer distribution tilts in favor of the entrepreneur, the exposure to a higher wage bill limits the realization of potential profits.

Intuitively as entrepreneurs hire more workers, they increase the scale of their operation. The bigger the size of the firm, the greater its expected sales, but also the greater the entrepreneur's potential loss. Conditional on their level of safe assets, entrepreneurs will choose a level of employment consistent with their expected sales and potential losses. What follows is a characterization of the model's equilibrium as well as its steady state dynamics.

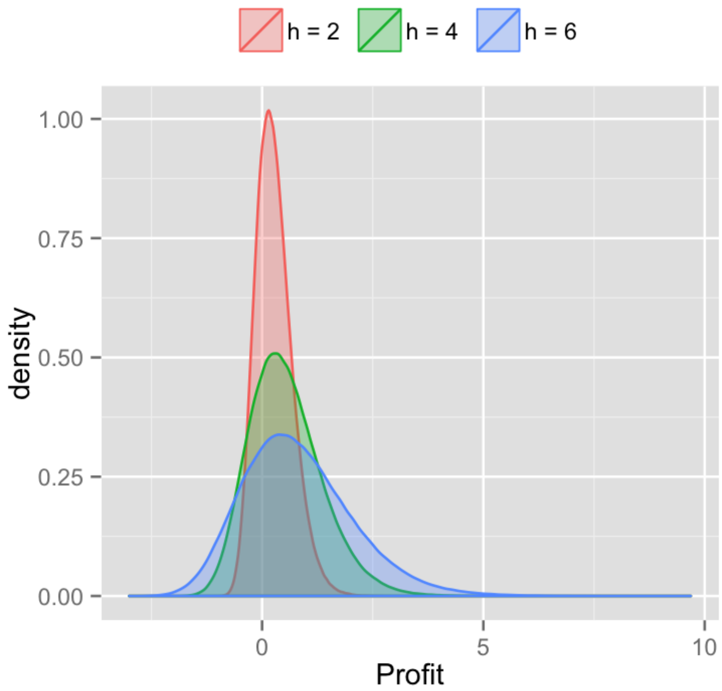

Figure 10: Profit distribution simulations

3.6 Equilibrium

Households do not face idiosyncratic risk and maximize lifetime utility by choosing c_t, h_t and b_{t+1} for all $t = 0, 1, 2, ...$ Let $S_t = \{B_t, \theta_t\}$ represent the aggregate states of the economy at time t, namely the extent of credit conditions in the economy and the aggregate level of wealth. For convenience, the household's recursive problem is restated below:

$$\begin{aligned} V(S,b) &= \max_{c,h,b'} \left\{ c - \gamma \frac{h^{1+\tau}}{1+\tau} + \beta E_\theta V(S',b') \right\} \\ s.t. \quad &: \quad wh + \frac{b'}{R} \geq c + b \\ &: \quad \Omega \geq \frac{b'}{R} \end{aligned}$$

Their policies satisfy the first order conditions:

$$\begin{aligned} \gamma h_t^\tau &= w_t \\ u_c(c_t, h_t) &\geq \beta R_t E_t u_c(c_{t+1}, h_{t+1}) \end{aligned}$$

where the last condition will be satisfied with equality if the borrowing constraint is binding. Similarly, the recursive problem for firm i at time t could be written as:

$$V(S,b^i) \quad = \quad \max_{h^i} E_n \left\{ \max_{b^{i'}} \left[\ln \left(y^i h^i + b^i - wh^i - \frac{b^{i'}}{R} \right) + \beta E_\theta V(S', b^{i'}) \right] \right\}$$

where E_θ refers to the expectation of θ_{t+1} conditional on θ_t and E_n refers to the unconditional expectation over all potential realizations of n_t^i. This difference resides in that n_t^i does not exhibit any serial correlation, while θ_t does. Given the above, we can define a recursive competitive equilibrium as follows:

Definition 1 *A Recursive Competitive Equilibrium consists of the following functions:*

(a) A value function $V_E(B,\theta,b^i)$ and decision rules $c^i(B,\theta,b^i)$, $h^i(B,\theta,b^i)$ and $b^{i'}(B,\theta,b^i)$ for the entrepreneur

(b) A value function $V_H(B,\theta,b)$ and decision rules $c(B,\theta,b)$, $h(B,\theta,b)$, and $b'(B,\theta,b)$ for the household

(c) Price functions w(B, θ) and R(B, θ)

(d) A perceived law of motion for the aggregate state $S' = \Phi(S) = \Phi(B,\theta)$

such that:

(i) Given c) and d), a) solves the entrepreneur's optimization problem

(ii) Given c) and d), b) solves the household's optimization problem

(iii) All markets clear:

$\int c_t^i \ dF(i) + c_t^H = Y_t$ *(Goods market)*

$\int \phi_t b_t^i \ dF(i) = \alpha h_t^\tau$ *(Labor market)*

$\int (b_t^i - \frac{b_{t+1}^i}{R_t}) \ dF(i) = b_t + \frac{b_{t+1}}{R_t}$ *(Financial markets)*

(iv) Perceptions about the aggregate states are correct

Implicit in the equilibrium's definition is the presence of the end of period Walrasian market described above. In turn, agents may seek to maximize lifetime consumption, independently on the number of consumption goods they originally acquired.

Given the equilibrium definition above, the model's solution is detailed below. Since the choice of labor h_t^i is made before the realization of the matching shock n_t^i, but the saving decision is made after its observation, it will be convenient to define the entrepreneur's wealth after production has taken place as:

$$a_t^i = b_t^i + (\theta_t n_t^i - w_t^i) h_t^i$$

Rewriting per-worker sales y_t^i in terms of θ_t and n_t^i, and defining ϕ_t as the value that satisfies $E_n \left[\frac{\theta_t n_t^i - w_t}{(\theta_t n_t^i - w_t)\phi_t + 1} \right] = 0$, then as shown by Quadrini (2017) the entrepreneur's policy functions will take the form:

$$\begin{aligned} h_t^i &= \phi_t b_t^i \\ c_t^i &= (1-\beta) a_t^i \\ b_{t+1}^i &= \beta R_t a_t^i \end{aligned}$$

Note that the demand for labor will be linear in the entrepreneur's wealth b_t^i. The factor of proportionality is time-varying, but common to all firms. In turn, the aggregate demand for labor can be obtained as:

$$H_t = \phi_t \int_{i \in N} b_t^i = \phi_t B_t$$

where B_t denotes the average, per-capita level of wealth, and the factor of proportionality ϕ_t will depend negatively on the equilibrium wage rate. In turn, this implies that the aggregate demand of labor will depend negatively on the wage rate (as in any Walrasian model), but positively on the economy's level of riskless assets. For individual producers these assets represent a firm's financial net worth. The corresponding theoretical proof can be found in Appendix 7.1.

The above is a unique feature of the model which sheds some light on the relationship between labor demand and the financial soundness of firms. When businesses' net worth suffer (as it does during contractions), the demand for labor declines inducing a lower equilibrium output and employment. This happens not as a result of firms lacking the resources to hire employees, or because the value of their collateral has plummeted and access to financing options are scarce. This occurs purely out of risk considerations: with a lower net worth entrepreneurs cannot properly insure against idiosyncratic shocks and seek to reduce their exposure by limiting their hiring. In other words, given the fact that entrepreneurs are risk averse and cannot hedge their hiring bets appropriately, they choose to behave conservatively and revise their production plans downwards. The opposite will happen in an expansion when a firm's net worth improves. This is a unique and a crucial feature of the model which will greatly affect the equilibrium dynamics as described in the next sections.

Another property worth mentioning is that an entrepreneur's consumption policy function is linear in wealth. This has two major implications. First, it implies that entrepreneurs will always consume (and save) a constant proportion of their end of period assets. As such, during expansions entrepreneurs will not only seek to consume more but also to increase their stock of savings which will allow them to increase future production. Second, it makes the problem extremely tractable as it allows for linear aggregation. Consequently, even when entrepreneurs might be heterogeneous in asset holdings, in order to understand the aggregate dynamics we only need to keep track of the average level of wealth B_t.

Lastly, in a stationary equilibrium, households will exhaust their credit capacity as long as $\beta R < 1$. Given that entrepreneurs are risk averse and face uninsurable idiosyncratic risks, they will constantly seek to self-insure. Their desire to smooth consumption would make them save and hold bonds even if $\beta R = 1$. Unfortunately for them, the supply of these assets is constrained by the borrowing limit of households. Being risk neutral and solely exposed to an aggregate shock, households need extra incentives to issue the risk free assets. In turn, in order to induce households to borrow the equilibrium interest must decline. As long as the interest rate is lower than the intertemporal discount rate, households will continue to increase their leverage until their borrowing limit binds setting the steady state interest rate lower than the intertemporal discount rate.

4 Calibration

The model parameters are calibrated to match a set of relevant moments from U.S. data. There are two sets of parameters. The first set of parameters is chosen externally without using model-generated data while the second set of parameters is determined jointly by minimizing the distance between the statistics from the model and the data.

The model period is a year, which corresponds to the data frequency obtained from Compustat. I set $\bar{y}$ to match the U.S. long-run, mean capacity utilization measures of approximately eighty percent as reported by the Federal Reserve. Following Whalen and Reichling (2017), I set $\tau = 0.4$, implying a labor elasticity of 2.5. This number is also in line with what is used and recommended by the U.S. Congressional Budget Office as in Reichling (2012). The persistence of the aggregate financial shock is estimated as an AR(1) process from the survey of senior loan officers available since the second quarter of 1990. Further, I use customer traffic data to estimate σ_n and set $\mu_n = 0$ since the model features constant returns to scale and consequently μ_n will only have a scaling effect on the economy. The rest of the parameters (β, γ, Ω) are calibrated to match the following steady state moments: U.S. long-run real interest rate of 3%, hours worked = 1/3, and the ratio of unsecured credit to disposable income at 0.41 as reported by Braxton et al. (2020) and Herkenhoff (2019). Table 2 below summarizes this information.

Table 2: Calibration Values

Parameter	Description	Value	Target/Source
β	Discount factor	0.952	Interest rate $r = 3\%$
γ	Disutility of labor	1.33	Hours worked = 1/3
Ω	Borrowing limit	0.117	Unsecured Credit/ Income = 0.41
τ	Inv. Frisch elasticity	0.40	CBO estimate (2012)
μ_n	Parameter of matching function	0.00	Consumer Traffic data
σ_n	Parameter of matching function	0.122	Consumer Traffic data
ρ_θ	Persistence of credit shock	0.859	FRB Senior Loan Officer Survey
σ_θ	Stdev of credit shock	0.008	FRB Senior Loan Officer Survey
$\bar{y}$	Maximum output per worker	0.902	FRB U.S. Capacity utilization rate

4.1 Customer Traffic data

Central to the study's analysis is an understanding of the distribution of customers served by firms in each period. I employ granular proprietary customer foot-traffic data from ShopperTrak to gain a better understanding on customer distribution[2]. The dataset contains proprietary consumer foot-traffic information for about 3,780 stores, all located within the United States. The data are annual and span four years, from 2010 to 2013, with an average of 70 billion visits recorded each year. The sample is geographically diversified, with all 50 U.S. states represented. Because of privacy considerations, the actual brands included in the sample were not disclosed; however, an anonymous numeric identifier makes it possible to track individual zip codes over time. Overall, the dataset forms a balanced panel with a total of 280 billion recorded visits. Table 3 reports the key summary statistics.

Table 3: ShopperTrak Data Moments

Statistic	Value
Consumer traffic cross-sectional dispersion	1.213
Consumer traffic cross-sectional skewness	-0.307
Consumer traffic cross-sectional kurtosis	1.257
Consumer traffic growth rate corr w/ cycle	0.407^{**}
Mean number of visits per store, per year	18.6 Mil
Mean number of visits per year	70 Bil
Total number of visits in sample	281 Bil
Unique ZIP-level units	378
Annual observations per ZIP level	4

Source: Own calculations based on consumer traffic data

Table 3 highlights substantial cross-sectional heterogeneity in consumer traffic across ZIP-level markets. The distribution of log customer traffic is mildly negatively skewed and has negative excess kurtosis, indicating somewhat thinner tails than a normal distribution in logs. At the same time, the underlying traffic data in levels are strictly positive, highly dispersed, and characterized by a long right tail driven by a relatively small number of very high-traffic ZIP-year observations. These empirical features make a lognormal specification a useful parsi-

[2]Founded in 1989 in Chicago and now part of Johnson Controls/Sensormatic retail solutions portfolio, ShopperTrak employed electronic traffic counters (ETC) to quantify and monitor customer movements inside as well as in and out stores. At the time the data was collected, the company had about fifty thousand ETC devices installed only in North America and about seventy thousand worldwide.

monious approximation: it imposes positive support, captures substantial cross-sectional heterogeneity, and maps an approximately bell-shaped distribution in logs into a right-skewed distribution in levels.

Furthermore, Table 3 shows that consumer traffic growth is positively and significantly correlated with the business cycle, indicating that store visits tend to rise during expansions and weaken during downturns. This procyclical pattern is not limited to the ShopperTrak micro data. It is also visible in aggregate consumer-traffic measures, including diffusion indices produced by the International Council of Shopping Centers (ICSC), which track changes in shopping-center traffic over time. As shown in Figure 18 in the Appendix, the ICSC series declines sharply during recessionary periods, most notably during the 2007–09 recession, and subsequently recovers as macroeconomic conditions improve. Taken together, the micro-level ShopperTrak moments and the aggregate ICSC evidence suggest that customer traffic is closely tied to the business cycle, supporting the model's treatment of store traffic as a cyclical object.

Finally, to discipline the mapping from observed store traffic to the model's traffic-per-worker object, I first estimate the elasticity of employment with respect to customer traffic using annual changes in aggregate employment and the ShopperTrak traffic index. I then use this estimate to proxy for the cross-sectional staffing response across stores. The estimated elasticity is close to 0.9, implying that staffing moves in a nearly proportional manner with customer traffic over time. Under this mapping, the cross-sectional dispersion of log traffic per worker is approximately $(1-\alpha)$ times the cross-sectional dispersion of log store traffic. Using the customer-traffic dispersion estimate of 1.213, this implies a traffic-per-worker dispersion of about 0.12. Thus, the value used in the model is disciplined jointly by the co-movement of aggregate employment and customer traffic and by the observed cross-sectional heterogeneity in store traffic.

5 Results

This section examines the quantitative implications of the model. It first evaluates the model's ability to reproduce key features of the Compustat data, with particular emphasis on the cross-sectional distribution of firm sales and employment. It then studies the dynamic response of the model economy to an aggregate credit shock. Finally, it decomposes the response into a first-moment, or level, effect and an endogenous uncertainty effect, thereby quantifying the contribution of time-varying firm-level uncertainty to aggregate fluctuations.

5.1 General Results

Table 4 below reports some fundamental simulation results. The basic strategy was to calibrate the model utilizing steady state moments and then validating the framework with non-targeted ones at the business cycle frequency[3]. By construction, the calibration disciplines the key parameters using the average degree of cross-sectional heterogeneity in firm outcomes and the distribution of customer traffic. Overall the calibrated model performs reasonably well in reproducing the targeted steady-state features of the data.

Table 4: Targeted Moments

Moment	Data	Model
Steady State interest rate	0.030	0.033
Hours worked	0.333	0.334
Unsecured debt/ Income	0.410	0.411

In addition, the model can successfully replicate several important features of the non-targeted moments. Table 5 summarizes some of these results. The model generates substantial dispersion in both employment and sales, consistent with the empirical evidence that firms differ markedly in size and demand conditions. Furthermore, the model also successfully replicates the higher level of cross-sectional dispersion of sales relative to employment, with sales being about 10 percent more dispersed than employment. Quantitatively, however, the model somewhat overstates the cross-sectional dispersion of both employment and sales, which partially reflects the parsimonious nature of the model, in which both firm employment

[3]Since the framework has a closed form solution, I am implicitly assuming that the steady state moments are equal to the model's ergodic mean; something which in principle is only assured for linearized models. In turn, I perform a consistency check which can be found in the Appendix 7.6.

and sales are driven by a common underlying source of idiosyncratic demand heterogeneity.

Table 5: Non-targeted Moments

Moment	Data	Model
Cross-sectional dispersion of employment	2.373	3.112
Cross-sectional dispersion of sales	2.619	3.435
Sales growth rate dispersion corr. w/ cycle	-0.194	-0.170
Emp. growth rate dispersion corr. w/ cycle	0.156	0.019
Employment's cross-sectional skewness	-0.053	-0.358
Sales' cross-sectional skewness	-0.319	-0.358
Employment's cross-sectional kurtosis	2.386	2.667
Sales' cross-sectional kurtosis	2.894	2.667

The correlation between sales-growth dispersion and the aggregate cycle is negative in both the data and the model, indicating that the model reproduces the countercyclical nature of sales uncertainty. In the data, this correlation is -0.194, while the model generates -0.170. This empirical regularity has also been documented by other researchers using different data sets. For example, Bachmann and Bayer (2013) report similar results for Germany using USTAN data. The model also matches the qualitative procyclicality of employment-growth dispersion, although its quantitative fit is weaker than for sales-growth dispersion. The model generates the right business cycle co-movement as an improvement in credit conditions induces firms to raise their sales forecasts and consequently increase their hiring. This procyclicality in employment is consistent with the behavior of customer traffic, which also moves positively with the business cycle in both granular consumer-traffic data and aggregate measures as in Figure 18 of the appendix.

Finally, the model does reasonably well in capturing the higher moments of the log-transformed cross-sectional distributions. Both in the data and in the model, employment and sales display mildly negative skewness, indicating modest asymmetry after the log transformation and trimming. The model somewhat overstates the negative skewness of employment, but remains close to the empirical value for sales. In terms of kurtosis, the empirical distributions are close to, or slightly below, the Gaussian benchmark of 3, suggesting that the log-transformed and trimmed data do not exhibit pronounced fat tails. The model reproduces this feature by generating kurtosis below 3 for both employment and sales. Overall, the model succeeds in reproducing the central economic mechanism of interest–substantial cross-sectional heterogeneity and countercyclical demand uncertainty, especially as reflected in the dispersion of sales growth.

5.2 Model Dynamics

In addition to the model's steady-state properties, its dynamic features are studied through impulse response functions. Figure 11 plots the response of aggregate output, employment, sales per worker, and the real wage to a one percent positive innovation to aggregate credit conditions, θ_t. On impact, both output and employment increase. Aggregate output rises by about 0.8 percent, while employment rises by roughly 0.6 percent. The remaining difference is accounted for by an increase in sales per worker, which rises by approximately 0.18 percent. Thus, the expansion is driven primarily by higher labor demand, but also by a modest increase in worker productivity.

The increase in labor demand is driven primarily by two effects. First, a higher θ_t raises the expected demand for goods faced by firms by increasing the expected level of sales per worker. As a result, the marginal worker generates more revenue in expectation, leading firms to increase labor demand. This is the standard first-moment, or level, effect and, as shown in the next section, it accounts for the dominant share of the employment response. Second, a higher level of aggregate credit reduces the demand uncertainty faced by entrepreneurs by compressing the cross-sectional dispersion of sales per worker. This makes the return to hiring an additional worker less risky and further increases firms' desired employment. The top-left panel of Figure 12 shows that demand uncertainty falls by close to eight percent on impact as higher credit availability compresses the cross-sectional distribution of sales per worker.

Overall the shock raises the expected payoff from hiring and therefore increasing firms' desired employment per unit of assets. The resulting expansion in labor demand puts upward pressure on the real wage, which increases by about 0.25 percent on impact and then gradually returns toward its steady-state level. Similarly, capacity utilization rates rise sharply as firms update their production plans to meet the expected growth in demand for consumption goods.

A positive credit shock raises demand but also pushes more firms against the per-worker capacity ceiling. In the simulation, the share of firms operating at maximum capacity rises by roughly 2.2 percentage points on impact, a significant response relative to the one percent innovation in θ_t that generated it. This capacity margin helps explain why the response of sales per worker is smaller than the aggregate shock.

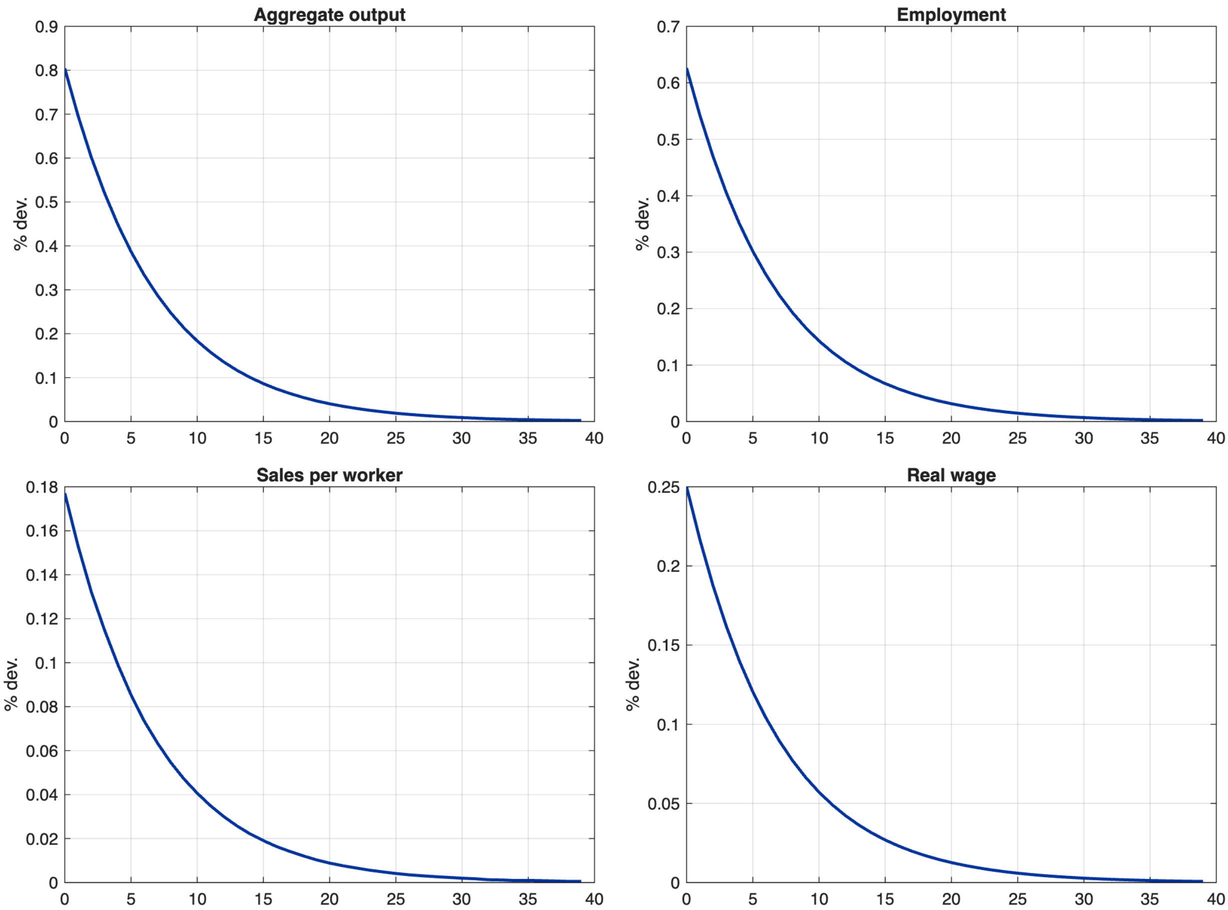

Figure 11: Impulse responses for a 1% shock to θ_t

The reduction in uncertainty induced by the positive credit shock also affects interest rates and the accumulation of safe assets. As higher credit availability reduces the effective uncertainty faced by entrepreneurs about future sales, their precautionary demand for safe assets declines. This lower demand puts downward pressure on bond prices and raises the equilibrium interest rate. Consistent with this mechanism, the model implies a decline in aggregate entrepreneurial assets and an increase in the interest rate on impact as seen in Figure 12.

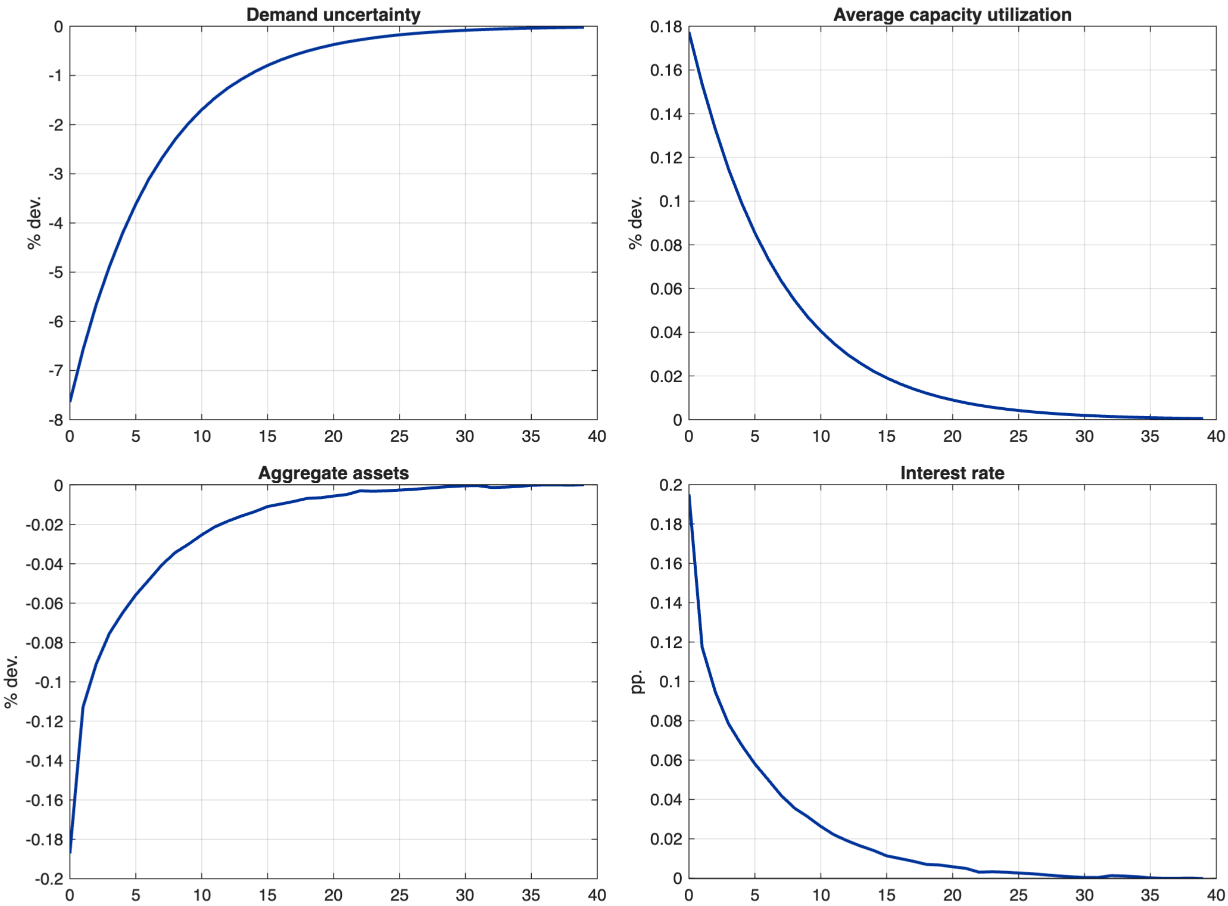

Figure 12: Impulse responses for a 1% shock to θ_t

5.3 Effect Decomposition

Section 5.2 illustrates how changes in aggregate credit conditions affect the economy through two distinct channels. The first is a standard first-moment, or level, effect: a higher value of θ_t raises expected sales per worker and therefore increases the expected marginal revenue from hiring. The second is a second-moment effect: because sales per worker are subject to a capacity constraint, changes in θ_t also alter the cross-sectional distribution of firm-level sales and employment. This endogenous change in dispersion amplifies the initial level effect of the aggregate shock by changing the riskiness of firms' hiring decisions. When credit conditions improve, the effective dispersion of sales per worker falls, making the payoff from hiring an additional worker less uncertain. Risk-averse entrepreneurs therefore increase labor demand by more than they would in response to a pure level shift in expected sales. The purpose of this section is to quantify how much of the aggregate response to a credit shock is due to this second channel.

To isolate the role of endogenous uncertainty, the baseline framework is compared with a counterfactual economy exposed to the same aggregate credit shock, but constructed so that the shock affects only the first moment of firm-level demand. In the baseline model, a change in θ_t alters both the mean and the cross-sectional dispersion of sales per worker because it shifts the effective capacity threshold faced by firms. The counterfactual economy removes this distributional effect while preserving the same response of average sales per worker:

$$y^i_{t,\text{level}} = A_t \min\{n^i_t, \bar{y}\} \tag{1}$$

where the scalar A_t is chosen each period so that average sales per worker in the counterfactual equals average sales per worker in the baseline economy under the same realization of θ_t:

$$A_t = \frac{E\left[\min\{\theta_t n^i_t, \bar{y}\}\right]}{E\left[\min\{n^i_t, \bar{y}\}\right]} \tag{2}$$

This construction preserves the level effect of the credit shock while shutting down the endogenous change in cross-sectional uncertainty generated by the moving capacity margin. Thus, any difference between the baseline and counterfactual impulse responses can be attributed to the endogenous uncertainty channel.

Figures 13 and 14 report the resulting decomposition. The level-only economy accounts for most of the expansion in aggregate activity, confirming that the standard first-moment channel is quantitatively important. On impact, aggregate output rises by about 0.61 percent in the counterfactual economy, compared with about 0.80 percent in the baseline economy. Thus, the endogenous uncertainty channel accounts for roughly one quarter of the total output response. The effect is even more pronounced for employment: the level-only economy generates an increase of about 0.43 percent, while the baseline model implies an increase of about 0.63 percent. Hence, approximately one third of the employment response is due to the decline in demand uncertainty. This amplification occurs because the positive credit shock not only raises expected sales per worker, but also compresses the cross-sectional dispersion of sales per worker, making hiring less risky for entrepreneurs.

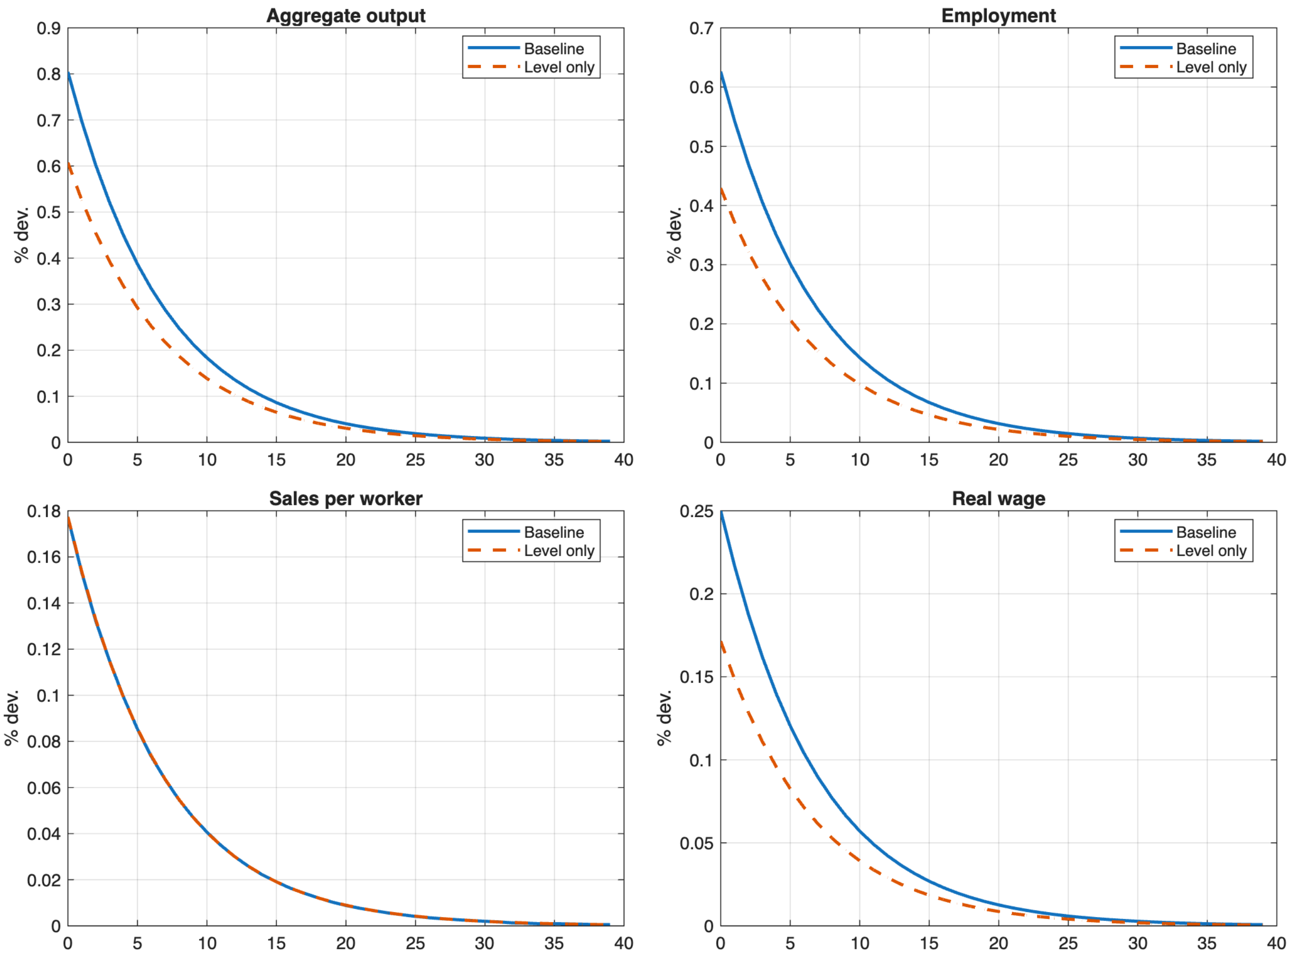

Figure 13: Effect decomposition for a 1% shock to aggregate credit

The decomposition also clarifies the behavior of prices and financial variables. Since employment expands more strongly in the baseline economy, the real wage rises by more than in the level-only counterfactual. In contrast, sales per worker behaves almost identically across the two economies, as intended by the construction of the counterfactual: the scalar A_t is chosen precisely to match the first-moment response of average sales per worker. The main difference therefore comes from the change in the distribution around that mean. This distributional effect is also visible in asset markets. In the baseline economy, the decline in demand uncertainty reduces entrepreneurs' precautionary demand for safe assets, causing aggregate entrepreneurial assets to fall and the equilibrium interest rate to rise. In the level-only economy, where this uncertainty channel is absent, assets instead rise slightly and the interest rate falls modestly. The gap between the two impulse responses therefore captures the financial-market counterpart of the endogenous uncertainty channel. Overall, the decomposition shows that while the level effect explains the majority of the response to a credit shock, endogenous uncertainty provides a sizable amplification mechanism, especially for employment, wages, asset demand, and interest rates.

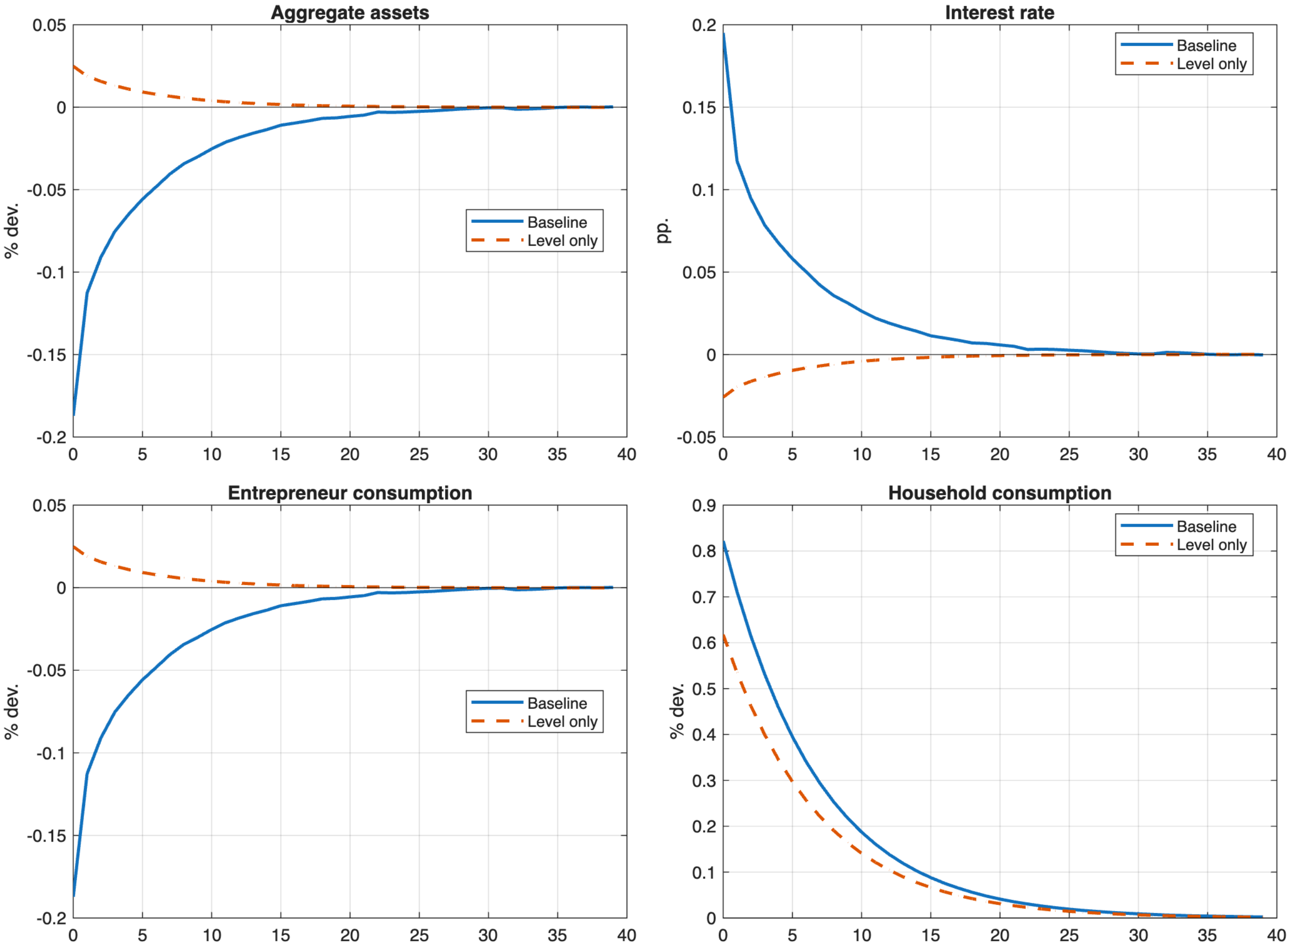

Figure 14: Effect decomposition for a 1% shock to aggregate credit

5.4 Sensitivity Analysis

This section explores the sensitivity of the results reported above to variations in two key parameters of the model: the elasticity of labor supply, governed by τ, and the per-worker capacity ceiling, $\bar{y}$. These parameters are especially relevant for the quantitative strength of the model's propagation mechanism. The labor supply elasticity determines how strongly wages and employment adjust in equilibrium following a change in labor demand. A more elastic labor supply allows employment to respond more strongly with a smaller increase in wages, whereas a less elastic labor supply shifts more of the adjustment into wages and dampens the employment response. The capacity ceiling $\bar{y}$ governs the extent to which firms are constrained in translating higher customer demand into sales per worker. When the ceiling binds more tightly, changes in aggregate credit conditions have stronger effects on the cross-sectional distribution of sales per worker and therefore on the endogenous uncertainty channel. By contrast, when $\bar{y}$ is sufficiently high, fewer firms operate near capacity and the model behaves more like a standard level-shock economy. The exercises below assess whether the main quantitative conclusions, particularly the amplification of employment and output through endogenous demand uncertainty, are robust to alternative values of these parameters.

5.4.1 The Role of Labor-Supply Elasticity

Intrinsically linked to the response of employment supply to variations in credit conditions, the parameter τ plays an important role in the overall dynamics of the model. Trade-off between fluctuations in the real wage and the labor supply will condition the response of output in the economy. Table 6 below explores the responses of the model to variations in the parameter governing the labor supply of households. In each case, the model was recalibrated utilizing the same targets specified in section 4, but each time with a particular value for τ. The tabulated results are the parameter values as well as the response (on impact) of employment and real wages to a one percent improvement in credit conditions measured as a percent deviation from their steady state values.

As expected, the response of employment to a positive credit shock becomes stronger as

Table 6: Model's sensitivity to τ

τ	γ	β	CR	Employment	Wage
0.2	1.106	0.952	0.117	1.23%	0.223%
0.4	1.330	0.952	0.117	0.63%	0.232%
0.5	1.472	0.952	0.117	0.50%	0.236%
0.7	1.804	0.952	0.117	0.36%	0.244%
1.0	2.447	0.952	0.117	0.25%	0.251%

labor supply becomes more elastic, that is, as τ declines. When τ is low, households are more willing to adjust hours in response to a given change in wages, so the expansion in labor demand is absorbed primarily through higher employment. When τ is high, labor supply is less elastic and the same shock generates a smaller increase in employment. In the recalibrated economies, however, the impact response of the real wage is very similar across values of τ. This reflects the labor-supply condition, $w_t = \gamma H_t^\tau$: the decline in the employment response as τ rises is offset by the larger elasticity of wages with respect to employment. Thus, the sensitivity exercise shows that τ primarily affects the quantity margin of labor adjustment, while the wage response remains quantitatively similar across calibrations.

5.4.2 The effects of Capacity Utilization

The degree to which productive capacity is used in the economy is an important factor governing its cross-sectional dynamics. In particular, when the steady-state share of firms operating at their per-worker capacity constraint is lower, a larger fraction of firms is able to expand production plans and translate a positive innovation to aggregate credit conditions into higher sales per worker, rather than immediately becoming constrained. This strengthens the first-moment, or level, effect of the shock relative to the baseline calibration. At the same time, because fewer firms are initially close to the constraint, the shock generates a smaller reduction in demand uncertainty than in the baseline economy. Thus, the response in the lower-hit-rate calibration is driven less by the compression of cross-sectional dispersion and more by the direct expansion in expected demand. The stronger increase in sales per worker raises the expected marginal revenue from hiring, leading firms to expand employment more aggressively. In turn, the greater the share of firms initially operating below full capacity, the greater the increase in labor demand and hence the greater the rise in real wages.

To illustrate this mechanism, this section performs the following counterfactual experiment. The model is recalibrated so that a counterfactual economy begins from a lower steady-state capacity hit rate than in the baseline calibration. Specifically, while the baseline economy is calibrated so that roughly 80 percent of firms operate at the per-worker sales ceiling $\bar{y}$, the counterfactual economy is calibrated to have a steady-state capacity hit rate of 70 percent. All other parameters are kept fixed. This exercise therefore isolates how the economy's response to the same aggregate credit shock changes when a larger share of firms initially has unused productive capacity. Figure 15 plots the dynamic response of the baseline and counterfactual economies to a one percent positive innovation to aggregate credit. The comparison highlights that when fewer firms are initially constrained, the shock generates a stronger level response: more firms are able to translate the increase in credit conditions into higher sales per worker, which raises labor demand and puts stronger upward pressure on real wages, even when the reduction in demand uncertainty is less pronounced.

On impact, aggregate output and employment rise by about fifty percent more than in the baseline calibration. The stronger output and employment responses reflect the larger first-moment effect of the aggregate shock. Since fewer firms are initially constrained by the per-worker sales capacity, more firms can translate the improvement in aggregate credit conditions into higher sales per worker. In this sense, the lower-hit-rate economy behaves more like an economy with unused productive slack: the shock raises the scale of activity for many firms before the capacity margin becomes binding. Consistent with this mechanism, sales per worker and real wages also rise more in the counterfactual economy than in the baseline. At the same time, the decline in demand uncertainty is somewhat smaller than in the baseline economy, since fewer firms are close enough to the capacity constraint for the shock to generate the same degree of compression in the cross-sectional distribution. This illustrates the sensitivity induced by the capacity margin: when more firms have slack, the level effect of the shock becomes more prominent, while the endogenous reduction in uncertainty, although still present, is somewhat attenuated. Average capacity utilization nevertheless increases more strongly, reflecting the fact that initially unconstrained firms move closer to their productive

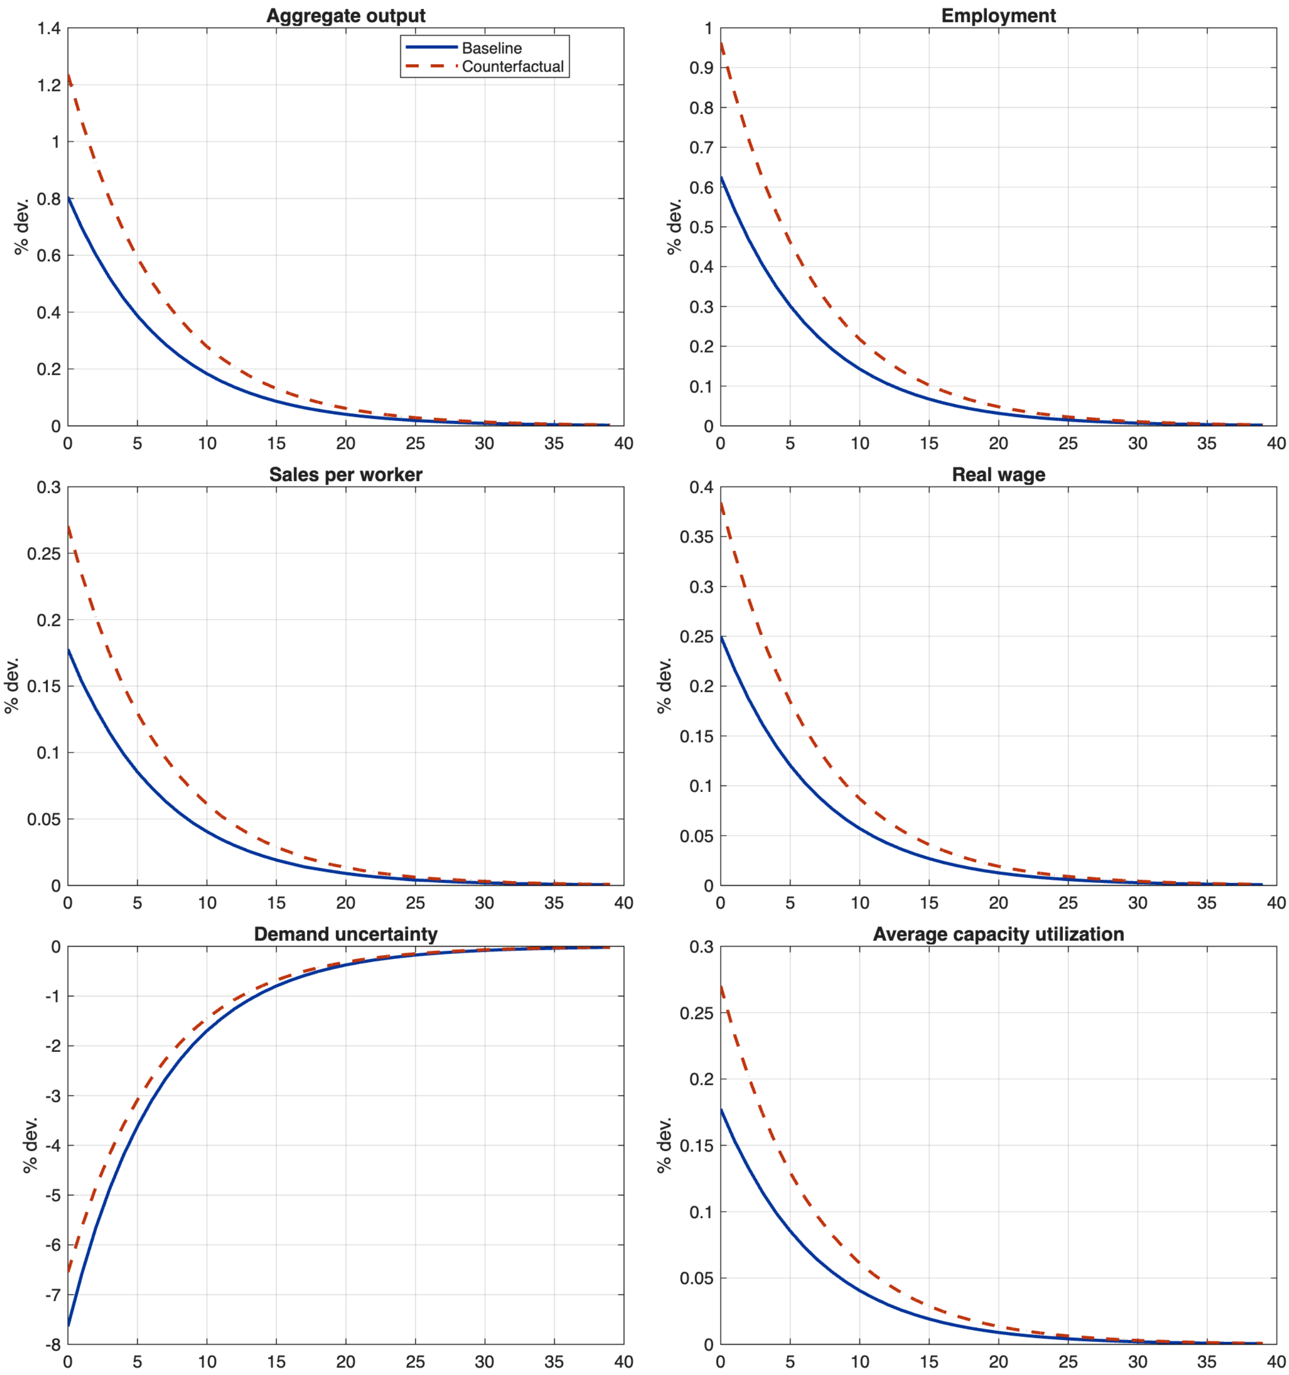

Figure 15: Effects of a 1% shock to θ_t under alternative capacity utilization levels.

capacity. Overall, the exercise confirms that the model's qualitative dynamics are robust, while showing that the quantitative strength of the expansion depends importantly on how much slack capacity firms have in steady state.

6 Conclusion

This study investigates the role of uncertainty fluctuations in shaping aggregate economic activity. In particular, it explores the hypothesis that changes in uncertainty need not be purely exogenous, but may instead arise endogenously from the current state of the economy. To study this mechanism, the paper develops a general equilibrium incomplete-markets framework with heterogeneous firms, idiosyncratic demand risk, and aggregate fluctuations in credit conditions. The key feature of the model is that expansions and contractions are initiated by changes in aggregate credit conditions, which affect not only the expected level of firms' demand but also the amount of uncertainty faced by individual producers. In this way, first-moment shocks generate endogenous second-moment effects.

The model generates realistic volatility in aggregate output and employment and provides a mechanism through which uncertainty can amplify the real effects of aggregate disturbances. In the baseline calibration, the endogenous uncertainty channel accounts for approximately twenty to thirty percent of the total response to a credit shock. Thus, fluctuations in uncertainty are not merely a separate source of aggregate volatility; they also operate as a propagation mechanism that strengthens the effects of more conventional level shocks. The model also predicts that uncertainty is countercyclical: it declines during expansions and rises during contractions. This implication is consistent with the empirical evidence for the United States, where standard measures of uncertainty tend to fall in booms and increase during recessions.

The analysis also shows that aggregate shocks have important effects on the cross-sectional distribution of firm outcomes. Changes in credit conditions alter not only average output, employment, and consumption, but also the dispersion of sales, employment, and firm-level risk. This highlights the importance of accounting for heterogeneity across producers and for the way in which individual risk exposure changes over the business cycle. Aggregate time series alone may therefore obscure an important part of the transmission mechanism, since the same aggregate shock can have different effects depending on the distribution of firms relative to their capacity constraints.

Overall, the results suggest that a full understanding of business cycles requires studying both aggregate dynamics and cross-sectional distributions. Theories of fluctuations should explain not only movements in mean output, employment, and consumption, but also why the distribution of firm behavior changes over the cycle and how those distributional changes feed back into aggregate outcomes. The interaction between aggregate credit conditions, firm-level uncertainty, and production decisions offers a quantitatively relevant channel for understanding the amplitude of business cycles and the dynamics of job creation and destruction.

7 Appendix

7.1 Omitted Theoretical Proofs

Proposition 1. Individual labor demand is linear in financial wealth (b_t^i), while consumption and savings are linear in total assets (a_t^i):

$$\begin{aligned} h_t^i &= \phi_t b_t^i \\ b_{t+1}^i &= R_t \beta a_t^i \\ c_t^i &= (1-\beta) a_t^i \end{aligned}$$

Proof Proposition 1.

The recursive formulation of the entrepreneur's problem presented in section 3.6 can also be written in terms of the information available to the agent at the time of making a decision. In turn, I define the following two stages or sub-problems:

Stage I:

$$\begin{aligned} V_t(\theta_t, B_t, b_t^i) &= \max_{h_t^i} E_{n_t} \hat{V}_t(\theta_t, B_t, a_t^i) \\ s.t. &: \quad a_t^i = (\theta_t n_t^i - w_t) h_t^i + b_t^i \end{aligned}$$

Stage II:

$$\begin{aligned} \hat{V}_t(\theta_t, B_t, a_t^i) &= \max_{c_t^i} [\ln c_t + \beta E_{\theta_{t+1}} V_{t+1}(\theta_{t+1}, B_{t+1}, b_{t+1}^i)] \\ s.t. &: \quad a_t^i \geq c_t^i + \frac{b_{t+1}^i}{R_t} \end{aligned}$$

where $E_{\theta_{t+1}}$ stands for the expectation of θ_{t+1} conditional on the realization of θ_t.

In stage I the entrepreneur chooses its labor inputs aware of the extent of credit conditions, yet uncertain about the level of demand that he will receive that period. In stage II, the entrepreneur observes the realization of n_t^i and allocates the end of period wealth between consumption and savings. The stage I first order condition is:

$$\frac{\delta V_t}{\delta h_t^i} \iff E_{n_t}\left[\frac{\delta \hat{V}_t}{\delta a_t^i}\frac{\delta a_t^i}{\delta h_t^i}\right] = 0$$

The envelope condition $\delta \hat{V}_t / \delta a_t^i = 1/c_t^i$ is derived and then used in the expression above to yield:

$$\frac{\delta V_t}{\delta h_t^i} \iff E_{n_t}\left[\frac{\theta_t n_t^i - w_t}{c_t^i}\right] = 0$$

The stage II first order condition is:

$$\frac{\delta \hat{V}_t}{\delta c_t^i} = 0 \iff \frac{1}{c_t^i} + \beta E_{n_t}\left[E_{\theta_{t+1}}\frac{\delta V_{t+1}}{\delta b_{t+1}^i}(-R_t)\right] = 0$$

Substituting the relevant envelope condition, and denoting E_t as conditional expectation given the information set at time t yields the following Euler equation:

$$\frac{1}{c_t^i} = \beta E_t R_t \left(\frac{1}{c_{t+1}^i}\right)$$

Next I prove Proposition 1 following a guess-and-verify approach. Begin by guessing the following policy functions:

$$h_t^i = \phi_t b_t^i \tag{3}$$

$$b_{t+1}^i = R_t \beta a_t^i \tag{4}$$

Replacing (4) in the stage II budget constraint

$$c_t^i = a_t^i - \frac{b_{t+1}^i}{R_t} \tag{5}$$

yields the policy function for consumption:

$$c_t^i = (1-\beta) a_t^i \tag{6}$$

From the Euler equation (FOC of stage II) we have that

$$\begin{aligned} \frac{1}{c_t^i} &= \beta R_t E_t \left(\frac{1}{c_{t+1}^i} \right) \\ &\Rightarrow \frac{1}{a_t^i} = \beta R_t E_t \left(\frac{1}{a_{t+1}^i} \right) \end{aligned} \tag{7}$$

Combining the definition of a_t^i and (3) yields

$$a_{t+1}^i = [(\theta_{t+1} n_{t+1}^i - w_{t+1})\phi_{t+1} + 1] b_{t+1}^i \tag{8}$$

which implies that (7) can be written as:

$$\begin{aligned} \frac{1}{a_t^i} &= \left(\frac{\beta R_t}{b_{t+1}^i} \right) E_t \left(\frac{1}{1 + (\theta_{t+1} n_{t+1}^i - w_{t+1})\phi_{t+1}} \right) \\ &\Rightarrow 1 = E_t \left(\frac{1}{1 + (\theta_{t+1} n_{t+1}^i - w_{t+1})\phi_{t+1}} \right) \end{aligned} \tag{9}$$

For the proof to be complete I need to verify that (9) satisfies the problem's FOCs:

$$E_t \left[\frac{\theta_t n_t^i - w_t}{(\theta_t n_t^i - w_t)\phi_t + 1} \right] = 0 \tag{10}$$

In turn, from (9)

$$E_t \left[\frac{1}{1 + (\theta_t n_t^i - w_t)\phi_t} \right] - 1 = 0 \tag{11}$$

$$\Rightarrow E_t \left[\frac{1 - 1 - (\theta_t n_t^i - w_t)\phi}{1 + (\theta_t n_t^i - w_t)\phi_t} \right] = 0$$

$$\Rightarrow (-\phi) E_t \left[\frac{\theta_t n_t^i - w_t}{1 + (\theta_t n_t^i - w_t)\phi_t} \right] = 0$$

$$\Rightarrow E_t \left[\frac{\theta_t n_t^i - w_t}{(\theta_t n_t^i - w_t)\phi_t + 1} \right] = 0 \tag{12}$$

which satisfies (10).

7.2 Aggregate Measures

For this economy, aggregate real income will equal the profits of the entrepreneurs and the labor income of the representative household. In turn:

$$
\begin{aligned}
Y_t &= \int (y_t^i - w_t) h_t^i \; dF(i) + w_t h_t \\
&= \int y_t^i h_t^i \; dF(i) - \int w_t h_t^i \; dF(i) + w_t h_t \\
&= \int y_t^i h_t^i \; dF(i)
\end{aligned} \tag{13}
$$

In terms of real consumption:

$$
\begin{aligned}
c_t^{e,i} dF(i) &= (y_t^i - w_t) h_t^i + b_t^i - \frac{b_{t+1}^i}{R_t} \\
&\Rightarrow \int c_t^{e,i} \; dF(i) = \int (y_t^i - w_t) h_t^i \; dF(i) + \int b_t^i \; dF(i) - \int \frac{b_{t+1}^i}{R_t} \; dF(i)
\end{aligned}
$$

This implies that the aggregate consumption of entrepreneurs can be written as:

$$C_t^E = Y_t - w_t h_t + b_t^e - \frac{b_{t+1}^e}{R_t}$$

and aggregate consumption of the households as

$$C_t^H = w_t h_t + \frac{b_{t+1}^H}{R_t} - b_t^H$$

Hence total consumption in the economy would be equal to:

$$
\begin{aligned}
C_t^E + C_t^H &= Y_t - w_t h_t + b_t^e - \frac{b_{t+1}^e}{R_t} + w_t h_t + \frac{b_{t+1}^H}{R_t} - b_t^H \\
&= Y_t + b_t^e - \frac{b_{t+1}^e}{R_t} + \frac{b_{t+1}^H}{R_t} - b_t^H \\
&= Y_t
\end{aligned}
$$

which is the total income/production described by expression 13.

7.3 Alternatives Measures of Uncertainty

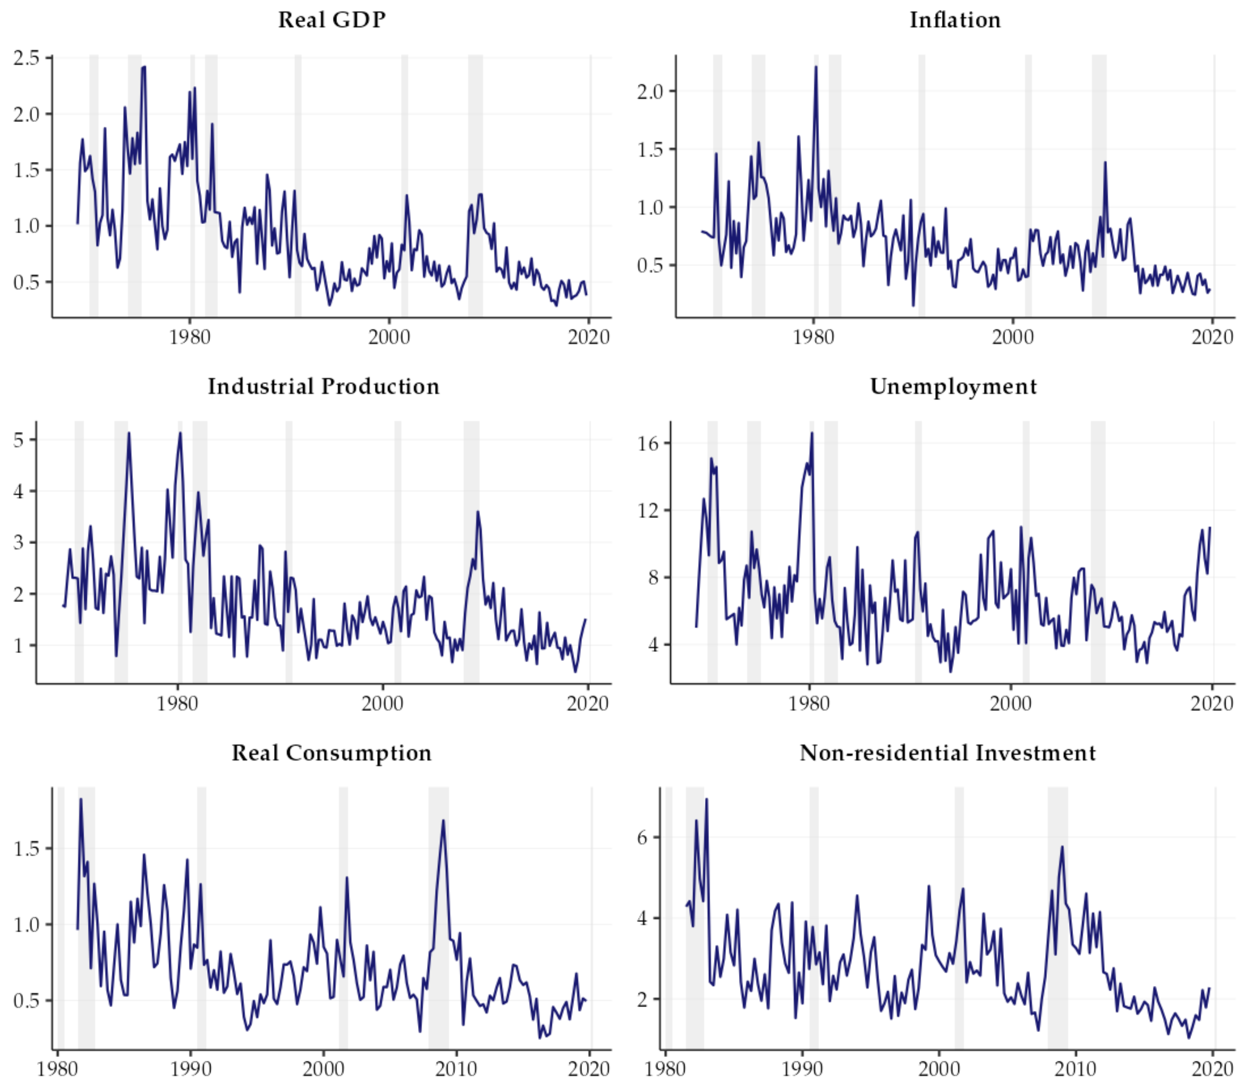

Figure 16: Disagreement amongst professional forecasters. The figure above plots the cross-sectional dispersion in private sector forecasts over the business cycle. The data comes from the Federal Reserve Bank of Philadelphia's survey of professional forecasters from 1968Q4 - 2014Q3 for the first four variables and 1981Q3 - 2014Q3 for the remaining two. Beginning from top left we have the forecasts for *Real GDP*, the *Price Deflator*, *Industrial Production*, the *Unemployment rate*, *Real Consumption* and *Non-residential fixed investment*. In times of higher uncertainty forecasts become less precise and dispersion amongst predictions increases. Not surprisingly, recessions tend to be periods of greatest disagreement amongst forecasters.

7.4 Evidence of Corporate Lending

In the framework introduced in Section 3 resources would, in equilibrium, flow from the entrepreneurs to the household sector. At first this result might seem like an odd feature of the model. However, in the U.S., the private corporate sector has been a net lender since the beginning of the 2000s as seen in figure 17. The only exception to date has been the year 2008 at the height of the Great Recession, when the financial assets held by most corporations dropped in value.

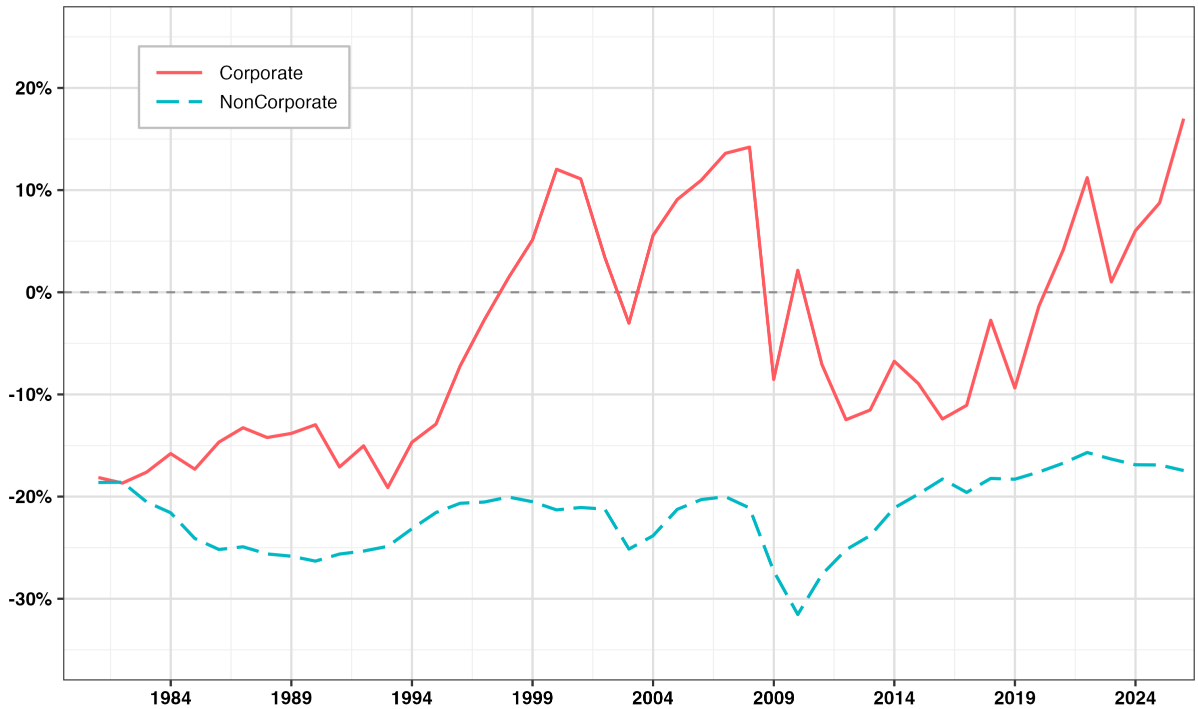

Figure 17: Net Financial Assets in the nonfinancial business sector as a percentage of total nonfinancial assets (1980-2025). **Source:** Federal Reserve Z.1 Financial Accounts.

Interestingly the reversal from net borrower to lender has so far only occurred in the U.S. Corporate sector, and not in the Noncorporate one. The evidence reported in the figure above shows that a large fraction of the business sector is self-financing and no longer dependent on outside sources. And even when the aggregate figures may mask some firm level heterogeneity, they do paint a general picture of the evolution of the overall trend across time.

7.5 Customer Foot-traffic over the Cycle

Figure 18 plots a monthly customer-traffic diffusion index from the International Council of Shopping Centers (ICSC), based on monthly survey responses from shopping-center executives. The index is centered around 50, with readings above 50 indicating a general positive momentum in the number of customers visiting shopping centers and readings below 50 indicating weakening traffic conditions. The series suggests that consumer traffic is strongly procyclical: it falls during recessions and recovers as economic conditions improve. This pattern is especially pronounced during the 2007–09 recession, when the index dropped sharply and remained below 50 for an extended period before gradually recovering. Overall, even when there is only enough data to cover two U.S. recessions, both the 2001 and 2008 downturns appear clearly visible and the figure provides additional evidence that shopping-center foot traffic moves closely with the business cycle and deteriorates substantially during periods of macroeconomic stress.

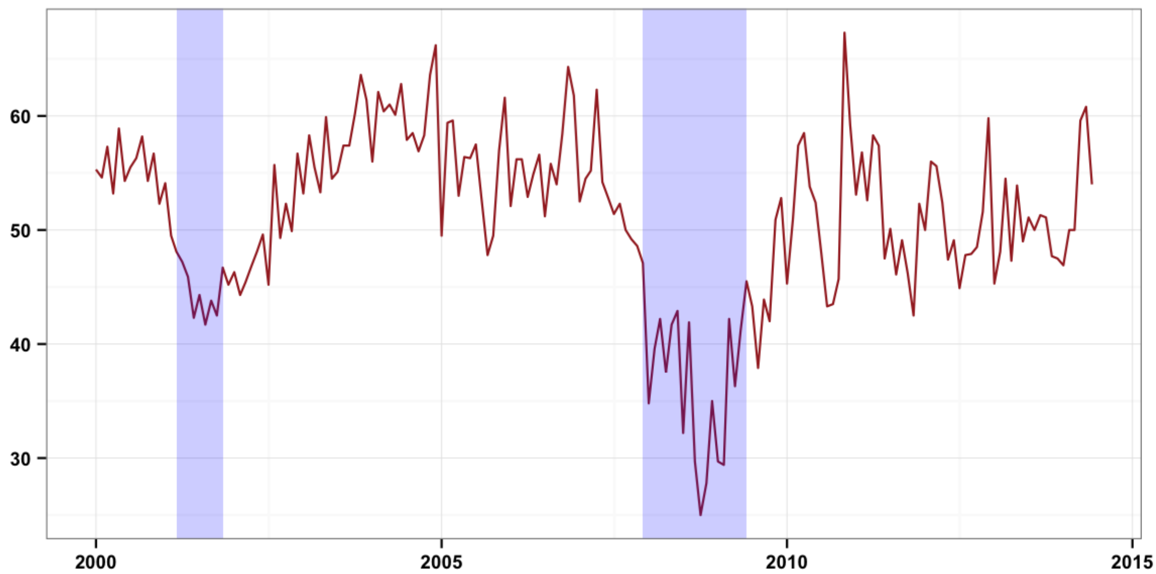

Figure 18: Consumer Traffic and Business Cycle
Source: Own calculations based on ICSC data

7.6 Validation Exercise

The calibration uses steady-state model moments as counterparts to the long-run moments of the stochastic economy. Since the model is nonlinear and features a capacity constraint, the equality between steady-state moments and ergodic averages is not mechanical. In particular, aggregate credit shocks shift both the mean level of sales per worker and the effective capacity threshold faced by firms, so Jensen effects could in principle make the long-run stochastic moments differ from the corresponding steady-state objects. To evaluate the quantitative importance of this issue, the baseline calibration is subjected to a simulation-based validation exercise.

The validation compares two simulated economies. In the first economy, aggregate credit conditions are fixed at their steady-state value, $\theta_t = 1$, in every period. This economy preserves the cross-sectional idiosyncratic demand risk faced by firms and therefore delivers the stationary counterpart to the steady-state moments used in the calibration. In the second economy, θ_t follows the calibrated stochastic process used in the baseline model. Both economies are simulated using the same equilibrium conditions, the same idiosyncratic customer-demand distribution, and the same baseline parameter values. Each economy is simulated for 10,000 periods with 200,000 firms per period, and the first 1,500 periods are discarded as burn-in. The reported ergodic moments are time averages over the post-burn-in sample.

The results show that the steady-state moments are very close to the ergodic averages of the stochastic economy. The differences in output, employment, wages, assets, debt-income ratios, sales per worker, and capacity utilization are all small in absolute value. The same is true for the cross-sectional moments: employment and sales dispersion differ by less than one-thousandth of a log point, while the differences in skewness and kurtosis are also quantitatively minor. Thus, although the nonlinear structure of the model means that equality between steady-state and ergodic moments is not guaranteed analytically, the simulation indicates that the approximation is accurate for the baseline calibration. This supports the use of steady-state moments as empirical counterparts in the calibration exercise.

Table 7: Validation of Steady-State and Ergodic Moments

Moment	Steady state	Ergodic mean	Difference
Mean output	0.322	0.322	-0.000
Mean employment	0.362	0.362	-0.000
Mean real wage	0.886	0.886	-0.000
Mean net interest rate (%)	3.102	3.104	0.002
Mean aggregate assets	0.116	0.116	-0.000
Mean debt/income	0.365	0.365	0.000
Mean sales per worker	0.890	0.890	-0.000
Capacity hit rate	0.801	0.801	-0.000
Employment dispersion	3.116	3.117	0.000
Sales dispersion	3.417	3.417	0.000
Employment skewness	0.071	0.072	0.001
Sales skewness	0.071	0.072	0.001
Employment kurtosis	1.863	1.862	-0.001
Sales kurtosis	1.864	1.863	-0.001

Notes: The steady-state economy fixes $\theta_t = 1$.

The ergodic mean is computed from the baseline stochastic economy after burn-in.

References

Andrew B Abel. Optimal investment under uncertainty. *The American Economic Review*, pages 228–233, 1983.

Andrew B Abel and Janice C Eberly. Optimal investment with costly reversibility. *The Review of Economic Studies*, 63(4):581–593, 1996.

S Rao Aiyagari. Uninsured idiosyncratic risk and aggregate saving. *The Quarterly Journal of Economics*, pages 659–684, 1994.

Iván Alfaro, Nicholas Bloom, and Xiaoji Lin. The finance uncertainty multiplier. *Journal of Political Economy*, 132(2):577–615, February 2024. doi: 10.1086/726230. URL `https://doi.org/10.1086/726230`.

David Altig, José María Barrero, Nicholas Bloom, Steven J. Davis, Brent Meyer, and Nicholas Parker. Surveying business uncertainty. *Journal of Econometrics*, 231(1):282–303, 2022. doi: 10.1016/j.jeconom.2020.03.021. URL `https://doi.org/10.1016/j.jeconom.2020.03.021`.

George-Marios Angeletos. Uninsured idiosyncratic investment risk and aggregate saving. *Review of Economic Dynamics*, 10(1):1–30, 2007.

George-Marios Angeletos and Laurent-Emmanuel Calvet. Idiosyncratic production risk, growth and the business cycle. *Journal of Monetary Economics*, 53(6):1095–1115, 2006.

George-Marios Angeletos and Vasia Panousi. Financial integration, entrepreneurial risk and global dynamics. *Journal of Economic Theory*, 146(3):863–896, 2011.

Cristina Arellano. Default risk and income fluctuations in emerging economies. *The American Economic Review*, pages 690–712, 2008.

Cristina Arellano, Yan Bai, and Patrick J. Kehoe. Financial frictions and fluctuations in volatility. *Journal of Political Economy*, 127(5):2049–2103, 2019. doi: 10.1086/701792. URL `https://doi.org/10.1086/701792`.

Yavuz Arslan, Aslihan Atabek, Timur Hulagu, and Saygin Sahinoz. Expectation errors, uncertainty and economic activity. *Oxford Economic Papers*, 67(3):634–660, 2015. doi: 10.1093/oep/gpv003. URL https://doi.org/10.1093/oep/gpv003.

Rüdiger Bachmann and Christian Bayer. 'wait-and-see'business cycles? *Journal of Monetary Economics*, 60(6):704–719, 2013.

Rüdiger Bachmann and Christian Bayer. Investment dispersion and the business cycle. *The American Economic Review*, 104(4):1392–1416, 2014.

Rudiger Bachmann and Giuseppe Moscarini. Business cycles and endogenous uncertainty. *Yale University mimeo*, 2011.

Rüdiger Bachmann, Steffen Elstner, and Eric R Sims. Uncertainty and economic activity: Evidence from business survey data. *American Economic Journal: Macroeconomics*, 5(2):217–249, 2013.

Yan Bai, José-Víctor Ríos-Rull, and Kjetil Storesletten. Demand shocks as technology shocks. *The Review of Economic Studies*, 2025. doi: 10.1093/restud/rdaf045. URL https://doi.org/10.1093/restud/rdaf045.

Scott R. Baker, Nicholas Bloom, and Steven J. Davis. Measuring economic policy uncertainty. *The Quarterly Journal of Economics*, 131(4):1593–1636, 2016. doi: 10.1093/qje/qjw024. URL https://doi.org/10.1093/qje/qjw024.

Scott R. Baker, Nicholas Bloom, and Stephen J. Terry. Using disasters to estimate the impact of uncertainty. *The Review of Economic Studies*, 91(2):720–747, 2024. doi: 10.1093/restud/rdad036. URL https://doi.org/10.1093/restud/rdad036.

Ravi Bansal and Amir Yaron. Risks for the long run: A potential resolution of asset pricing puzzles. *The Journal of Finance*, 59(4):1481–1509, 2004.

Susanto Basu and Brent Bundick. Uncertainty shocks in a model of effective demand. *Econometrica*, 85(3):937–958, 2017. doi: 10.3982/ECTA13960. URL https://doi.org/10.3982/ECTA13960.

Geert Bekaert, Marie Hoerova, and Marco Lo Duca. Risk, uncertainty and monetary policy. *Journal of Monetary Economics*, 60(7):771–788, 2013.

Truman Bewley. The permanent income hypothesis: A theoretical formulation. *Journal of Economic Theory*, 16(2):252–292, 1977.

Nicholas Bloom. The impact of uncertainty shocks. *Econometrica*, 77(3):623–685, 2009.

Nicholas Bloom. Fluctuations in uncertainty. *Journal of Economic Perspectives*, 28(2):153–176, 2014. doi: 10.1257/jep.28.2.153. URL `https://doi.org/10.1257/jep.28.2.153`.

Nicholas Bloom, Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten, and Stephen J. Terry. Really uncertain business cycles. *Econometrica*, 86(3):1031–1065, 2018. doi: 10.3982/ECTA10927. URL `https://doi.org/10.3982/ECTA10927`.

Nicholas Bloom, Steven J. Davis, Lucia S. Foster, Scott W. Ohlmacher, and Itay Saporta-Eksten. Investment and subjective uncertainty. NBER Working Paper 30654, National Bureau of Economic Research, Cambridge, MA, November 2022. URL `https://doi.org/10.3386/w30654`.

Nick Bloom, Stephen Bond, and John Van Reenen. Uncertainty and investment dynamics. *The Review of Economic Studies*, 74(2):391–415, 2007.

Benjamin Born and Johannes Pfeifer. Policy risk and the business cycle. *Journal of Monetary Economics*, 2014.

J. Carter Braxton, Kyle F. Herkenhoff, and Gordon M. Phillips. Can the unemployed borrow? implications for public insurance. NBER Working Paper 27026, National Bureau of Economic Research, Cambridge, MA, April 2020. URL `http://www.nber.org/papers/w27026`.

Giovanni Caggiano, Efrem Castelnuovo, and Nicolas Groshenny. Uncertainty shocks and unemployment dynamics in us recessions. *Journal of Monetary Economics*, 2014.

Dario Caldara, Cristina Fuentes-Albero, Simon Gilchrist, and Egon Zakrajšek. The macroeconomic impact of financial and uncertainty shocks. *European Economic Review*, 88:185–

207, 2016. doi: 10.1016/j.euroecorev.2016.02.020. URL https://doi.org/10.1016/j.euroecorev.2016.02.020.

Ambrogio Cesa-Bianchi, M. Hashem Pesaran, and Alessandro Rebucci. Uncertainty and economic activity: A multicountry perspective. *The Review of Financial Studies*, 33(8):3393–3445, 2020. doi: 10.1093/rfs/hhz098. URL https://doi.org/10.1093/rfs/hhz098.

Francisco Covas. Uninsured idiosyncratic production risk with borrowing constraints. *Journal of Economic Dynamics and Control*, 30(11):2167–2190, 2006.

Steven J Davis, R Jason Faberman, and John Haltiwanger. The flow approach to labor markets: New data sources and micro-macro links. *Journal of Economic Perspectives*, 20:3–26, 2006.

Ryan A. Decker, Pablo D'Erasmo, and Hernan J. Moscoso Boedo. Market exposure and endogenous firm volatility over the business cycle. *American Economic Journal: Macroeconomics*, 8(1):148–198, 2016. doi: 10.1257/mac.20130011. URL https://doi.org/10.1257/mac.20130011.

Stephanie Denis and Prakash Kannan. The impact of uncertainty shocks on the uk economy. *IMF Working Paper*, 2013.

Pablo N D'Erasmo and Hernan J Moscoso Boedo. Intangibles and endogenous firm volatility over the business cycle. *University of Virginia mimeo*, 2011.

Ian Dew-Becker and Stefano Giglio. Cross-sectional uncertainty and the business cycle: Evidence from 40 years of options data. *American Economic Journal: Macroeconomics*, 15(2):65–96, 2023. doi: 10.1257/mac.20210136. URL https://doi.org/10.1257/mac.20210136.

Pablo D. Fajgelbaum, Edouard Schaal, and Mathieu Taschereau-Dumouchel. Uncertainty traps. *The Quarterly Journal of Economics*, 132(4):1641–1692, 2017. doi: 10.1093/qje/qjx021. URL https://doi.org/10.1093/qje/qjx021.

J Fernández-Villaverde, P Guerrón-Quintana, J Rubio-Ramirez, and M Uribe. Risk matters: The real effects of stochastic volatility. *American Economic Review*, 2011.

Jesús Fernández-Villaverde, Pablo Guerrón-Quintana, Keith Kuester, and Juan F. Rubio-Ramírez. Fiscal volatility shocks and economic activity. *American Economic Review*, 105(11): 3352–3384, 2015. doi: 10.1257/aer.20121236. URL `https://doi.org/10.1257/aer.20121236`.

Ana Fostel and John Geanakoplos. Why does bad news increase volatility and decrease leverage? *Journal of Economic Theory*, 147(2):501–525, 2012.

Simon Gilchrist and John C Williams. Investment, capacity, and uncertainty: a putty–clay approach. *Review of Economic Dynamics*, 8(1):1–27, 2005.

Simon Gilchrist, Jae W Sim, and Egon Zakrajšek. Uncertainty, financial frictions, and investment dynamics. Technical report, NBER Working Paper No. 20038, 2014.

Luigi Guiso and Giuseppe Parigi. Investment and demand uncertainty. *Quarterly Journal of Economics*, pages 185–227, 1999.

Tarek A. Hassan, Stephan Hollander, Laurence van Lent, and Ahmed Tahoun. Firm-level political risk: Measurement and effects. *The Quarterly Journal of Economics*, 134(4):2135–2202, November 2019. doi: 10.1093/qje/qjz021. URL `https://doi.org/10.1093/qje/qjz021`.

Kyle F. Herkenhoff. The impact of consumer credit access on unemployment. *The Review of Economic Studies*, 86(6):2605–2642, 2019. doi: 10.1093/restud/rdz006. URL `https://doi.org/10.1093/restud/rdz006`.

Bernard Herskovic, Bryan Kelly, Hanno Lustig, and Stijn Van Nieuwerburgh. Firm volatility in granular networks. *Journal of Political Economy*, 128(11):4097–4162, 2020. doi: 10.1086/710345. URL `https://doi.org/10.1086/710345`.

Chris Higson, Sean Holly, and Paul Kattuman. The cross-sectional dynamics of the us business cycle: 1950-1999. *Journal of Economic Dynamics and Control*, 26(9):1539–1555, 2002.

Chris Higson, Sean Holly, Paul Kattuman, and Stylianos Platis. The business cycle, macroeconomic shocks and the cross-section: The growth of uk quoted companies. *Economica*, 71 (282):299–318, 2004.

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

Mark Huggett. The one-sector growth model with idiosyncratic shocks: Steady states and dynamics. *Journal of monetary economics*, 39(3):385–403, 1997.

Urban Jermann and Vincenzo Quadrini. Macroeconomic effects of financial shocks. *American Economic Review*, 102(1), 2012.

Kyle Jurado, Sydney C. Ludvigson, and Serena Ng. Measuring uncertainty. *American Economic Review*, 105(3):1177–1216, 2015. doi: 10.1257/aer.20131193. URL `https://doi.org/10.1257/aer.20131193`.

Matthias Kehrig. The cyclicality of productivity dispersion. *University of Texas at Austin mimeo*, 2011.

Miklos Koren and Silvana Tenreyro. Volatility and development. *The Quarterly Journal of Economics*, pages 243–287, 2007.

Per Krusell and Anthony A Smith. On the welfare effects of eliminating business cycles. *Review of Economic Dynamics*, 2(1):245–272, 1999.

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

John V Leahy and Toni M Whited. The effect of uncertainty on investment: Some stylized facts. *Journal of Money, Credit and Banking*, 28:64–83, 1996.

Sylvain Leduc and Zheng Liu. Uncertainty shocks are aggregate demand shocks. *Journal of Monetary Economics*, 82:20–35, 2016. doi: 10.1016/j.jmoneco.2016.07.002. URL `https://doi.org/10.1016/j.jmoneco.2016.07.002`.

Sydney C. Ludvigson, Sai Ma, and Serena Ng. Uncertainty and business cycles: Exogenous impulse or endogenous response? *American Economic Journal: Macroeconomics*, 13(4):369–410, 2021. doi: 10.1257/mac.20190171. URL `https://doi.org/10.1257/mac.20190171`.

Césaire A Meh and Vincenzo Quadrini. Endogenous market incompleteness with investment risks. *Journal of Economic Dynamics and Control*, 30(11):2143–2165, 2006.

Anna Orlik and Laura Veldkamp. Understanding uncertainty shocks and the role of black swans. *Journal of Economic Theory*, 222, 2024. doi: 10.1016/j.jet.2024.105905. URL `https://doi.org/10.1016/j.jet.2024.105905`.

Nicolas Petrosky-Nadeau and Etienne Wasmer. Macroeconomic dynamics in a model of goods, labor and credit market frictions. *Journal of Monetary Economics*, 72:97–113, 2015. doi: 10.1016/j.jmoneco.2015.01.006. URL `https://doi.org/10.1016/j.jmoneco.2015.01.006`.

Vincenzo Quadrini. The importance of entrepreneurship for wealth concentration and mobility. *Review of income and Wealth*, 45(1):1–19, 1999.

Vincenzo Quadrini. Entrepreneurship, saving, and social mobility. *Review of Economic Dynamics*, 3(1):1–40, 2000.

Vincenzo Quadrini. Bank liabilities channel. *Journal of Monetary Economics*, 89:25–44, 2017. doi: 10.1016/j.jmoneco.2017.03.006. URL `https://doi.org/10.1016/j.jmoneco.2017.03.006`.

Felix Reichling. How the supply of labor responds to changes in fiscal policy. Report, Congressional Budget Office, Washington, DC, October 2012. URL `https://www.cbo.gov/sites/default/files/cbofiles/attachments/10-25-2012-Labor_Supply_and_Fiscal_Policy.pdf`.

Chiara Scotti. Surprise and uncertainty indexes: Real-time aggregation of real-activity macro surprises. *Journal of Monetary Economics*, 82:1–19, 2016. doi: 10.1016/j.jmoneco.2016.06.002. URL `https://doi.org/10.1016/j.jmoneco.2016.06.002`.

LC Stein and EC Stone. The effect of uncertainty on investment, hiring, and r&d: Causal evidence from equity options. *Stanford University mimeo*, 2011.

Joseph Vavra. Inflation dynamics and time-varying volatility: New evidence and an S-s interpretation. *The Quarterly Journal of Economics*, 129(1):215–258, 2014. doi: 10.1093/qje/qjt027. URL `https://doi.org/10.1093/qje/qjt027`.

Charles J. Whalen and Felix Reichling. Estimates of the frisch elasticity of labor supply: A review. *Eastern Economic Journal*, 43:37–42, 2017. doi: 10.1057/eej.2015.23. URL `https://doi.org/10.1057/eej.2015.23`.