The Tail that Wags the Economy: Belief-Driven Business Cycles and Persistent Stagnation*

Julian Kozlowski† NYU

Laura Veldkamp‡ NYU Stern

Venky Venkateswaran§ NYU Stern

July 2, 2015

Abstract

The "great recession" was a deep downturn with long-lasting effects on credit markets, labor markets and output. We explore a simple explanation: This recession has been more persistent than others because it was perceived as an extremely unlikely event before 2007. Observing such an episode led all agents to re-assess macro risk, in particular, the probability of tail events. This change in beliefs endures long after the event itself has passed and through its effects on prices and choices, it produces long-lasting effects on investment, employment and output. To model this idea, we take a production economy and add agents who use standard econometrics tools to estimate the distribution of aggregate shocks. When they observe a new shock, they add that new piece of data to their data set and re-estimate the distribution from which it was drawn. Transitory shocks have persistent effects on beliefs because, once observed, the shocks stays forever in the agents' data set. We feed a time-series of data on actual macro shocks into our model, let our agents re-estimate the distribution from which the data is drawn each period, and show that our belief revision mechanism can explain the 13% downward shift in trend output.

JEL Classifications: E32

Keywords: Stagnation, tail risk, belief-driven business cycles

*Veldkamp and Venkateswaran gratefully acknowledge financial support from the NYU Stern Center for Global Economy and Business. We also thank Philippe Andrade, Christian Hellwig, Franck Portier, Robert Ulbricht for helpful comments and discussions.

†kozjuli@nyu.edu.

‡lveldkam@stern.nyu.edu.

§vvenkate@stern.nyu.edu.

1 Introduction

The "great recession" was a deep downturn with long-lasting effects on credit markets, labor markets and output. As Figure 1 shows, the financial crisis looks like a level shift in trend GDP. We explore a simple explanation: This recession has been more persistent than others because it was perceived as an extremely unlikely event. Observing these outcomes made us re-assess our estimates of macro risk. For example, in 2006, no one raised the possibility of financial panic. Today, the question of whether the financial crisis might repeat itself arises frequently and option prices continue to reflect heightened tail risk.

Importantly, these changes in beliefs can persist long after the event itself has passed. Since perceptions of tail risk affect prices and choices, this persistent change in beliefs has long-lasting output effects. To model this idea, we take a production economy and add agents who use standard econometric tools to estimate the distribution of aggregate shocks, in a non-parametric way. When they observe a new shock, they add that new piece of data to their data set and re-estimate the distribution from which it was drawn. Transitory shocks have persistent effects on beliefs because, once observed, the shocks stays forever in the agents' data set. We feed a time-series of data on actual macro shocks for the US into our model, let our agents re-estimate the distribution from which the data is drawn each period, and show that our belief revision mechanism can explain the 13% downward shift in trend output.

Although some post-war recessions have been mild and others deep, all of them had a distinct tough, followed by a sharp rebound toward trend outcome. None produced the kind of level-shift dynamic of the great recession. For standard macro models or even belief-driven business cycle models to produce persistent stagnation, they need to feed in sufficiently persistent negative shocks.[1] Not only is that approach at odds with the transitory 2008-09 shocks described by most data series, but it also provides no insight about why this recession looked so different. We propose a new type of belief-driven business cycle model where persistence is endogenous and state-dependent. The key difference is that our agents learn about the distribution of shocks, instead of a hidden or future state. We argue that fluctuations are persistent not because agents fear they are still in a "bad state," but rather, because the experience of new extreme events permanently changes their assessment of risk. This view is quantitatively successful, supported by data, and consistent with popular narratives. In Figure 2, the skew index, which measures tail risk from option prices, shows a rise in the financial crisis, with no subsequent decline. Furthermore, the financial sector's stagnation narrative emphasizes continued precautions against such risk. As Citibank (2014) writes, "All signs indicate that the recovery ... will be a slow one. Precautionary behavior on the part of banks will likely persist

[1] See Backus et al. (2015) for a formal analysis of propagation in business cycle models with belief shocks.

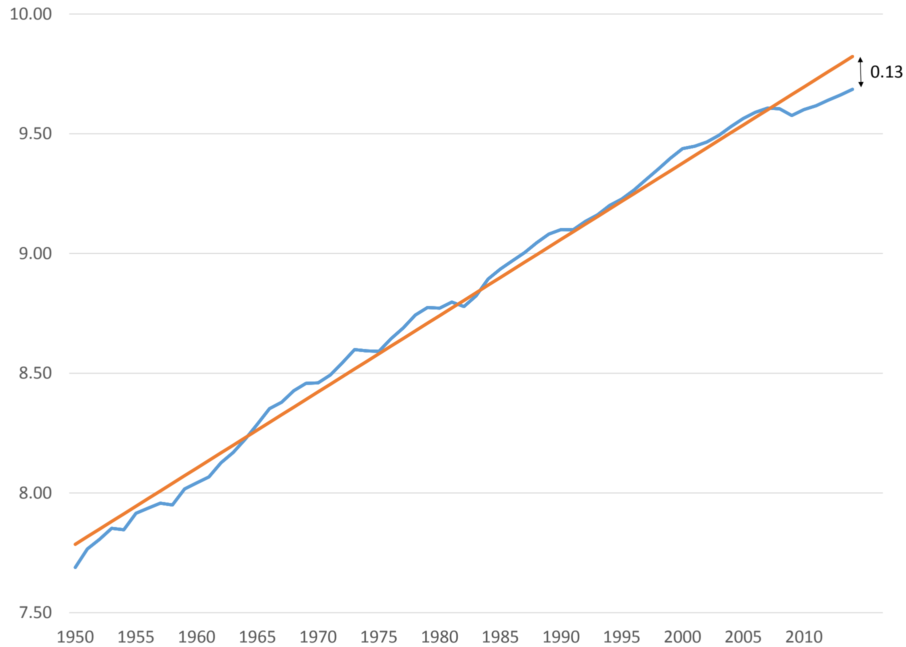

Figure 1: Real GDP in the U.S. and its trend, in logs. Blue line is a linear trend that fits data from 1950-2007. In 2014, real GDP was 0.13 below trend.

over the coming years." Following the 2008 crisis, not only banks, but also households exhibit far greater aversion to leverage. In our model, the symptom of higher perceived risk is such an aversion to lending. When borrowers and lenders believe tail risk, and thus the risk of costly default, is higher, they borrow and lend less.

Obviously, no one truly knows the distribution of economic shocks. That is a useful fiction economists use to discipline beliefs. Assuming that agents learn from data using standard econometric tools imposes just as much discipline. In addition, our approach has the advantage that it can produce persistent effects from transitory shocks; it can explain why measures of disaster risk fluctuate even in periods when no disaster is observed; and it is simple to execute. Our paper shows that this mechanism can generate persistent stagnation in an otherwise standard neoclassical setting. It also demonstrates how this econometric approach to belief updating ties beliefs firmly to observable data, and can be easily combined with sophisticated, quantitative macro models.

Not only does exploring econometric belief formation help quantitatively, it also provides new economic intuition. Investigating the process of statistical learning provides a new perspective on what it means for a price or quantity to be information insensitive. Typically, authors

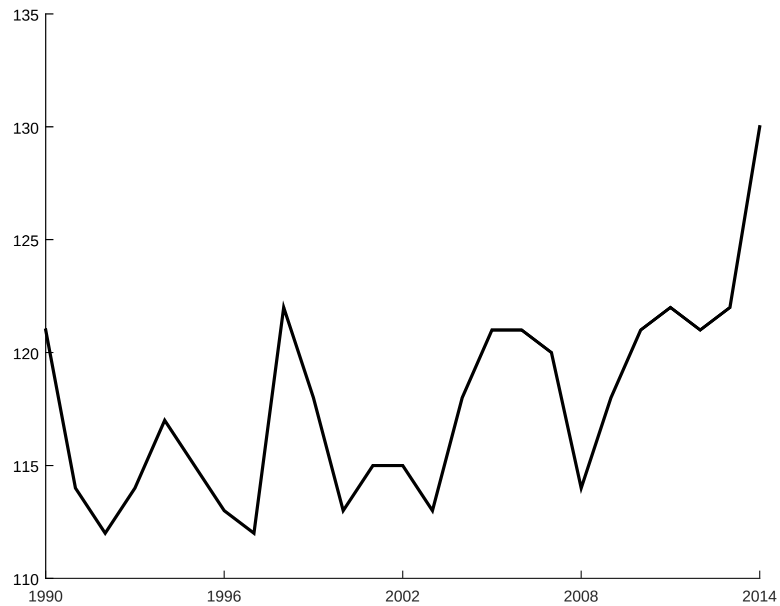

Figure 2: A 2-standard-deviation tail risk index (SKEW), constructed using option prices by the Chicago Board of Options. 1990M1:2014M3.

claim that debt is information-insensitive because its payoffs are flat throughout most of the state space. The one region where debt payoffs vary is in left-tail states that correspond to bankruptcy. Our learning mechanism reveals that beliefs in some regions of the state space are more sensitive to new information than others. In particular, beliefs about the probability of tail events are sensitive to new data because data in those regions is scarce. When data is scarce, new information strongly influences beliefs. Thus, debt is an instrument that is very sensitive to beliefs exactly in a region of the state space where beliefs are very sensitive to new information. Because of this effect, firms' use of debt makes the economy more information-sensitive. Similarly, curvature in preferences makes output more sensitive to left-tail events and raises information sensitivity as well. To illustrate this effect, we model firms that finance investment by issuing debt and show how a leveraged economy, where many firms have substantial debt, held by risk-averse households, is more sensitive to new observations.

Section 2 describes our quantitative, belief-driven business cycle model with defaultable debt. The model features a continuum of firms that produce output with capital and labor. These firms do not know the true distribution of aggregate shocks. Following many papers in the financial crisis literature, we use capital quality shocks (shocks to the effective capital depreciation rate) as the driving force of the model because this series has large, but transitory fluctuations in the crisis.[2] Each period, agents observe a new shock realization, add it to

[2]Some of the most well-known papers that have adopted this shock include Gertler and Karadi (2011), Brunnermeier and Sannikov (2014), and Gourio (2012).

their data set, and use a normal kernel-density estimator to re-estimate the shock distribution. Firms also experience idiosyncratic profit shocks. Because a continuum of idiosyncratic shocks is observed every period, the idiosyncratic shock distribution is perfectly observed and is therefore assumed to be common knowledge. Before observing either shock, each firm chooses its labor and capital investment. Wages and capital investment can be financed with debt, which yields a tax advantage to the firm, but also subjects it to bankruptcy costs if it defaults. We calibrate model parameters to match average leverage, including operating leverage, investment and default rates. The cost of issuing debt (the credit spread or risk premium) depends on the probability of default, which in turn, depends on the probability of adverse aggregate shocks. Thus, when the probability of a left tail event rises, the credit spread rises, debt issuance and real investment fall, and output declines.

The reason shocks have persistent effects is that once observed, a shock is in the agents' data set forever. The direct effect of the shock may pass quickly. But the observation of that event permanently alters the estimated probability of future events. The permanence reflects the martingale property of beliefs: If an agent estimates probabilities that will predictably rise or fall in the future, they should revise their estimate today. Efficient estimates of fixed objects in a model are always martingales, meaning that all innovations are permanent. Of course, if agents observe a sequence of moderate or positive shocks following the negative shock, the effect of the negative shock will diminish over time. But if the data observed in the future is consistent with the estimated distribution – mostly small shocks, with a few negative outliers – the left tail of the distribution will persist forever. Of course, permanent innovations is an extreme form of persistence and not necessary for us to explain the persistence of recent recession. Even if agents were to discount past data, the transitory observation still has persistent belief effects. Such discounting of past data would also offer one explanation for why the infinite history of the world economy (albeit with varying data quality) has not taught us the true distribution of macro shocks.

Section 5 quantifies the effect of our belief revisions. To do this, we construct a time-series of capital quality shocks using historical data on replacement and market value of the non-financial capital stock from the Flow of Funds report. Following the large negative shock to capital in 2008-'09, our agents increase their probability estimates of similar shocks. This increase in tail risk triggers a cumulative drop in capital of 15-20% and in output of about 12%, with almost no rebound to trend. These predicted effects of the financial crises, are similar in magnitude those in the data. Hall (2014) estimates that the U.S. capital stock and U.S. real GDP are each 13% lower than they would be if the economy had continued to grow at its pre-crisis rate of trend growth. Our model, which is driven by temporary decline capital quality and is calibrated to pre-crisis default and leverage data, matches the size of the output drop, without being calibrated to any GDP data. Hall also argues that the depressed rate of business capital formation was the single largest contributor to the persistent depressed output (often referred to as secular stagnation) in the post-crisis period. We use our model output to do the same decomposition as Hall does. We show that the symptoms of secular stagnation, as seen in changes in investment and labor between 2008 and 2015, resemble those in the data.

Next, we decompose this effect. When we turn off the belief-revision mechanism and endow agents with the distribution implied by the full data set, the initial impact of the capital shock is similar, but all macro aggregates immediately start to rebound. Similarly, we turn off risk aversion, debt financing, and revisions in tail risk. We find that each of these accounts for one-half to one-third of the long-run output effect. Finally, in order to show how persistence can vary, we use our model to predict the response to a much smaller economic shock. Specifically, we explore the predicted consequences of the shock that triggered the 2001 recession, which was about 1 standard deviation away from its sample mean. The model predicts fluctuations that are both smaller and less persistent. This exercise shows how, even if the martingale property implies that all belief revisions are permanent, cyclical persistence differs. When the permanent effect on beliefs is large relative to the transitory direct effect of the shock, the business cycle is more persistent. Such large belief revisions arise primarily when the shock is in a region of the state space where data is scarce. In other words, the 2001 recession was not very persistent because it was like many another recessions. In contrast, events that are unlike previously-observed events produce the most persistent output fluctuations.

Taken together, our results suggest that the recovery from the great recession may have has been slow simply because we learned that financial crises are still possible in the U.S. and this new knowledge permanently changed our assessment of macroeconomic risk.

Comparison to the literature The production side of our model builds on existing work that traces out the macro consequences of an exogenous increase in disaster risk, e.g. Gourio (2012) and is similar in many respects to the existing theories of shocks to beliefs that drive business cycles[3]. Our approach based on re-estimating the distribution of shocks offers two key advantages. First, our belief shocks are not exogenous. Without discipline on the possible time-series of beliefs, many macroeconomic outcomes are rationalizable. Our agents' beliefs are the outcome of a standard kernel-density estimation using actual macroeconomic data. The second advantage is that beliefs about fixed distributions are martingales, while beliefs about time-varying states are only persistent to the extent that one assumes the states or shocks

[3]Papers on news driven business cycles include papers on news shocks, such as, Beaudry and Portier (2004), Lorenzoni (2009), Veldkamp and Wolfers (2007), papers on uncertainty shocks, such as Jaimovich and Rebelo (2006), Bloom et al. (2014), Nimark (2014) and papers on higher-order belief shocks, such as Angeletos and La'O (2013) or Huo and Takayama (2015).

are persistent. Our mechanism delivers persistent effects of transitory shock and offers an explanation for why some shocks are more persistent than others. This can help us understand why many recessions have rapid recoveries and yet, some do not.

A small number of uncertainty-based theories of business cycles also deliver persistent effects from transitory shocks. In Straub and Ulbricht (2013) and Van Nieuwerburgh and Veldkamp (2006), a negative shock to output raises uncertainty, which feeds back to lower output, which in turn creates more uncertainty. To get even more persistence, Fajgelbaum et al. (2014) combine this mechanism with an irreversible investment cost, a combination which can generate multiple steady-state investment levels. Getting large persistent fluctuations from this kind of hidden state model is a challenge for two reasons. One hurdle is that states that are far apart are easy to distinguish. So uncertainty about which state one is in will be short-lived, even with noisy data. So, hidden states must be relatively close together, which makes the economic impact of state uncertainty small. The second hurdle is that a 5% drop in output is a large economic fluctuation. But a 5% drop in signal precision has small effects. So these models need to build in mechanisms to amplify changes on both margins.

These uncertainty-based explanations leave two questions unanswered. First, why were credit markets hardest hit and most persistently impaired after the crisis? Second, why did the depressed level of economic activity continue long after the VIX and other measures of uncertainty had recovered? Our theory relies more on tail risk. Like uncertainty, tail risk is a moment of the perceived distribution of outcomes. But the value of debt is particularly sensitive to this moment. The skew index data in Figure 2 reveal that tail risk has lingered, making it a better candidate for explaining continued depressed output.

Our belief formation process is similar to the parameter learning models by Johannes et al. (2015), Cogley and Sargent (2005) and Orlik and Veldkamp (2014) and is advocated by Hansen (2007). However, none of these papers has a production economy or considers persistent shocks to output. In Pintus and Suda (2015), parameter learning in a production economy amplifies shocks to leverage. While they feed in persistent leverage shocks and explore amplification, we feed in large shocks to capital and explore endogenous persistence. In addition, our non-parametric approach allows us to incorporate beliefs about tail risk.

Our model draws on many popular theories of the great recession, such as Gertler et al. (2010), Gertler and Karadi (2011), Brunnermeier and Sannikov (2014) and Gourio (2012). Moriera and Savov (2015) is similar to our model in that agents learn and it changes their demand for shadow banking (debt) assets. But their agents learn about a hidden two-state Markov process, which has persistence built in. While this literature has taught us an enormous amount about the mechanisms that triggered declines in lending and output in the financial crisis, it also hard-wired in persistent shocks. Our model aims to complement these theories by describing a simple mechanism that is compatible with many existing frameworks, is easy to implement, and delivers persistence. Rather than substituting for these existing narratives about the mechanics of financial crisis, our belief-formation mechanism adds another layer to the story, by explaining why some shocks deliver more persistent responses than others.

Finally, our paper contributes to the literature on secular stagnation. Eggertsson and Mehrotra (2014) claim that secular stagnation prevails because of negative real interest rates and low effective demand. While this may well be true, we offer a simple alternative. The key ingredient that transforms transitory business cycle shocks into persistent fluctuations is allowing agents to estimate the shock distribution in real time, just like an econometrician would. The specifics of our production economy, our debt-financed firms, and the use of capital quality shocks all function to make the persistent changes in beliefs large enough to deliver a quantitatively plausible theory of secular stagnation. But, more broadly, our model describes a belief-formation mechanism that can add persistent effects of extreme shocks in many frameworks.

2 Model

We explore the quantitative implications of parameter learning for persistence in a standard business cycle framework widely used in recent work on the financial crisis and the Great Recession. We begin by laying out a general model along the lines of Gertler and Karadi (2011) and Gourio (2012). Firms finance investment and payroll expenses using a combination of debt and equity financing and are subject to aggregate and idiosyncratic shocks. Our main innovation is to introduce real-time model estimation - specifically, agents in the model use the observed shocks and standard econometric tools (kernel density estimators) to estimate the shape of the distribution of aggregate shocks.

Preferences and technology: An infinite horizon, discrete time economy has a representative household, with preferences over consumption and labor supply, following

$$U_t = \left[(1-\beta) \left(C_t - \frac{\zeta L_t^{1+\gamma}}{1+\gamma} \right)^{1-\psi} + \beta E_t \left(U_{t+1}^{1-\eta} \right)^{\frac{1-\psi}{1-\eta}} \right]^{\frac{1}{1-\psi}} \tag{1}$$

where ψ is the inverse of the intertemporal elasticity of substitution, η indexes risk-aversion and γ is inversely related to the elasticity of labor supply.

The economy is also populated by a unit measure of firms, indexed by i and owned by the representative household. Firms produce output with capital and labor, according to a standard Cobb-Douglas production function $Ak_{it}^{\alpha} l_{it}^{1-\alpha}$, where A is total factor productivity (TFP), which is the same for all firms and constant over time. Firms are subject to an aggregate shock to capital quality ϕ_t. A firm that enters the period with capital $\hat{k}_{it}$ and is hit by a shock ϕ_t has effective capital $k_{it} = \phi_t \hat{k}_{it}$.

Our focus on capital quality shocks as the source of aggregate fluctuations is in the tradition of a number of recent papers on financial frictions, crises and the Great Recession - for example, Gertler et al. (2010), Gertler and Karadi (2011), Brunnermeier and Sannikov (2014) and Gourio (2012). These shocks work to *permanently* scale up or down the effective capital stock. However, such shocks by themselves are not enough to generate long-lived output responses. An adverse quality shock creates incentives to invest rapidly and return to a steady state level of capital. To deter this investment boom, Gertler and Karadi (2011), Gourio (2012) and others add persistence to the shock process - a bad shock not only wipes out a fraction of today's capital, but also makes it more likely that any investments today will be hit by bad shocks tomorrow. This persistence then spills over to aggregate outcomes, allowing them to generate long-lived output responses.

Importantly, we assume that the shock ϕ_t is i.i.d. The independence assumption ensures changes in beliefs are the only source persistence in the long run[4]. This distinguishes our results from those with exogenous persistence (i.e. assume that the ϕ_t are autocorrelated). The distribution of these shocks G is unknown to agents. Learning about G is the key novel feature and the focus of this paper. The following subsection will describe how agents use the observed history of ϕ_t to update their beliefs about G.

Firms are also subject to an idiosyncratic shock v_{it} . These shocks scale up and down the total resources available to each firm (before paying debt, equity or labor):

$$\Pi_{it} = v_{it} \left[A k_{it}^{\alpha} l_{it}^{1-\alpha} + (1-\delta) k_{it} \right] \tag{2}$$

where δ is the rate of capital depreciation. The shocks v_{it} are i.i.d across time and firms and are drawn from a known distribution[5], F. The mean of the idiosyncratic shock is normalized to be one: $\int v_{it} \quad di = 1$

Information Sets: The key innovation in the model is the assumption that agents must estimate the aggregate shock distribution G. Their common information set is $\mathcal{I}_t$, which includes all aggregate and idiosyncratic variables observed up to and including time-t. At each point in time, they use the empirical distribution of ϕ_t up to that point to construct an estimate $\hat{G}_t$ of

[4]In the short to medium run, consumption smoothing also introduces some persistence as the economy gradually recovers.

[5]This is a natural assumption - with a continuum of firms and a stationary shock process, firms can learn the complete distribution of any idiosyncratic shocks after one period.

the true distribution G. Formally, at every date t, agents construct the following kernel density estimator of the pdf g [6]:

$$\hat{g}_t(\phi) = \frac{1}{n_t \kappa} \sum_{s=0}^{n_t - 1} \Omega\left(\frac{\phi - \phi_{t-s}}{\kappa}\right)$$

where $\Omega(\cdot)$ is the standard normal density function, κ is the bandwidth parameter and n_t is the number of available observations of at date t. As new data arrives, agents update their estimates, generating a sequence of beliefs $\left\{\hat{G}_t\right\}$.

To keep our problem tractable, we follow most previous work on learning (Cogley and Sargent (2005), Piazzesi et al. (2015), Johannes et al. (2015)) in using anticipated utility (Kreps, 1998). At each date t, agents act as if the true distribution is $\hat{G}_t$. If we relax this assumption, it would not change how beliefs are updated. Nor would it affect the martingale properties of beliefs. Both are preference-independent. It would add a layer of uncertainty that would raise risk premia, as in Johannes et al. (2015). Since this extra uncertainty is not relevant to our main point about the persistence of real macro aggregates, and it adds considerable opacity and computational complexity, we suppress it.

Labor, credit markets and default: Firms hire labor in advance, i.e. before observing the realizations of aggregate and idiosyncratic shocks. Wages are non-contingent - in other words, workers are promised a non-contingent payment and face default risk.

Firms have access to a competitive non-contingent debt market, where lenders offer bond price (or equivalently, interest rate) schedules as a function of all relevant aggregate and idiosyncratic states, in the spirit of Eaton and Gersovitz (1981).

In order to characterize these schedules, we need to analyze the firm's default decision. A firm enters the period with an obligation, b_{it+1} to bondholders and a promise of $w_{it+1} l_{it+1}$ to its workers. The shocks are then realized and the firm (i.e. its shareholders) decide whether to repay their obligations or default. A firm that defaults makes no payments to equity holders. Formally, default is optimal for shareholders if

$$\Pi_{it+1} - b_{it+1} - w_{it+1} l_{it+1} + \Gamma_{t+1} < 0$$

where Γ_{t+1} is the present value of continued operations (we characterize this object later in this section - specifically, we will show that, since idiosyncratic shocks are iid, this is the same for all firms and, in equilibrium, equal to 0). Thus, the default decision is a function of the resources available to the firm (Π_{it+1}) and the *total* obligations of the firm to both bondholders

[6] In our numerical implementation, we fit a smooth density function to the empirical distribution. We also studied a handful of flexible parametric specifications, which yielded similar results.

and workers ($b_{it+1} + w_{it+1}l_{it+1} \equiv B_{it+1}$). The former is a function of the capital and labor choices, as well as the realizations of shocks. Let $r_{it+1} \in \{0,1\}$ denote the default policy of the firm.

In the event of default, the workers and bondholders take over the firm. The productive resources of a defaulting firm are sold to an identical new firm at a discounted price, equal to a fraction $\theta < 1$ of the value of the defaulting firm. The proceeds are distributed *pro-rata* among the creditors (both bondholders and unpaid workers). Note that the claims of both bondholders and workers have equal seniority[7].

Let $q\left(\hat{k}_{it+1}, l_{it+1}, B_{it+1}, S_t\right)$ denote the bond price schedule faced by a firm in period t. In other words, the firm receives $q(\cdot)$ in exchange for a promise to pay one unit of output at date $t+1$. Note that the bond price determination is made before the following period's capital quality shocks are known. Therefore, the price depends on the amount of capital invested $\hat{k}_{it+1}$, but it cannot be made contingent on the effective capital that will be available for production k_{it+1} or the profit shock v_{it+1}. The dependence on the other firm-level variables follows from our earlier discussion on the default decision. Formally,

$$q\left(\hat{k}_{it+1}, l_{it+1}, B_{it+1}, S_t\right) = \mathbb{E}_t M_{t+1}\left[r_{it+1} + (1 - r_{it+1})\frac{\theta\tilde{V}\left(\Pi_{it+1}, S_{t+1}\right)}{B_{it+1}}\right]. \tag{3}$$

where $\tilde{V}\left(\Pi_{it+1}, S_{t+1}\right)$ is the value of the assets of the firm (to be characterized later) and M_{t+1} is the stochastic discount factor of the representative household, which, given our Epstein-Zin specification takes the form

$$M_{t+1} = \left(\frac{dU_t}{dC_t}\right)^{-1}\frac{dU_t}{dC_{t+1}} = \beta\left[E_t\left(U_{t+1}^{1-\eta}\right)\right]^{\frac{\eta-\psi}{1-\eta}} U_{t+1}^{\psi-\eta}\left(\frac{u\left(C_{t+1}, L_{t+1}\right)}{u\left(C_t, L_t\right)}\right)^{-\psi} \tag{4}$$

Importantly, the bond price is a function of the aggregate state S_t, which includes the available history of aggregate shocks and outcomes. We will show later that S_t can be summarized by three objects - aggregate resources available, denoted Π_t, the labor input N_t, (which is chosen in advance, i.e. in $t-1$) and the estimated distribution $\hat{G}_t$.

Debt is assumed to carry a tax advantage, which creates incentives for firms to borrow. A firm which issues debt at price q_{it} and promises to repay b_{it+1} in the following period, receives a total date-t payment of $\chi q_{it} b_{it+1}$, where $\chi > 1$. This subsidy to debt issuance along with the cost of default introduces a trade-off in the firm's capital structure decision, breaking the Modigliani-Miller theorem[8].

[7]Note also that this means that default does not destroy resources - the penalty is purely private.

[8]The subsidy is assumed to be paid by a government that finances it through a lumpsum tax on the representative household.

For a firm that does not default, the dividend payout is its total available resources times output shock, minus its payments to debt and labor, minus the cost of building next period's capital stock (the undepreciated current capital stock is included in Π_{it}), plus the revenue earned from issuing new debt, including its tax subsidy:

$$d_{it} = \Pi_{it} - B_{it} - \hat{k}_{it+1} + \chi q_{it} b_{it+1} \tag{5}$$

Importantly, we do not restrict dividends to be positive, with negative dividends interpreted as (costless) raising of equity. Thus, firms are not financially constrained, ruling out another potential source of persistence.

Workers, who are also members of the representative family, evaluate their wage claims using the stochastic discount factor, M_{t+1}. This implies that the present value of a promise of wage w_{it+1} is given by

$$w_{it+1}\mathbb{E}_t M_{t+1}\left[r_{it+1} + (1 - r_{it+1})\frac{\theta\tilde{V}(\Pi_{it+1}, S_{t+1})}{B_{it+1}}\right] = w_{it+1}q_{it}$$

where the expectation is taken over aggregate and idiosyncratic shocks. From the household's problem, we can derive the following optimality condition for labor supply:

$$\begin{aligned} w_{it+1}q_{it}\frac{dU_t}{dC_t} &= \frac{dU_t}{dL_{t+1}} \\ w_{it+1}q_{it} &= \left(\frac{dU_t}{dC_t}\right)^{-1}\frac{dU_t}{dL_{t+1}} \\ &\equiv \mathcal{W}_t \end{aligned} \tag{6}$$

In other words, the expected value of wages, weighted by the economy-wide stochastic discount factor M_{t+1} is the same for all firms and is equal to the marginal rate of substitution of the representative household. The wage promise, w_{it+1}, must offer workers compensation for default risk. Since the risk is identical for bonds and wage payments, this risk adjustment involves simply multiplying the promised wage by the equilibrium bond price. In other words, the workers are essentially paid through bonds.

Timing, value functions and equilibrium: The timing of events in each period t is as follows:

1. Firms enter the period with a capital stock $\hat{k}_{it}$, labor l_{it}, outstanding debt b_{it}, and a wage obligation $w_{it}l_{it}$.

2. The aggregate capital quality shock ϕ_t and the firm-specific profit shock v_{it} are realized. Production takes place.

3. The firm decides whether to default or repay ($r_{it} \in \{0, 1\}$) its bond and labor claims.

4. The firm makes capital $\hat{k}_{it+1}$, debt b_{it+1} choices for the following period, along with wage/employment contracts w_{it+1} and l_{it+1}. Workers commit to next-period labor supply l_{it+1}. Note that all these choices are made concurrently.

Value functions: In recursive form, the problem of the firm is

$$V\left(\Pi_{it}, B_{it}, S_t\right) = \max\left[0, \max_{d_{it}, \hat{k}_{it+1}, b_{it+1}, w_{it+1}, l_{it+1}} d_{it} + \mathbb{E}_t M_{t+1} V\left(\Pi_{it+1}, B_{it+1}, S_{t+1}\right)\right] \tag{7}$$

subject to

$$\begin{aligned}
\text{Dividends:} \quad & d_{it} \ \leq \ \Pi_{it} - B_{it} - \hat{k}_{it+1} + \chi q_{it} b_{it+1} \\
\text{Discounted wages:} \quad & \mathcal{W}_t \ \leq \ w_{it+1} q\left(\hat{k}_{it+1},\, l_{it+1}, B_{it+1}, S_t\right) \\
\text{Future obligations:} \quad & B_{it+1} = b_{it+1} + w_{it+1} l_{it+1} \\
\text{Resources:} \quad & \Pi_{it+1} = v_{it+1}\left[A(\phi_{t+1}\hat{k}_{it+1})^{\alpha} l_{it+1}^{1-\alpha} + (1-\delta)\phi_{t+1}\hat{k}_{it+1}\right] \\
\text{Bond price:} \quad & q\left(\hat{k}_{it+1},\, l_{it+1}, B_{it+1}, S_t\right) = \mathbb{E}_t M_{t+1}\left[r_{it+1} + (1 - r_{it+1})\frac{\theta \tilde{V}_{it+1}}{B_{it+1}}\right]
\end{aligned}$$

The first max operator in (7) captures the firm's option to default if the value of the firm is negative. The expectation is taken over the idiosyncratic and aggregate shocks, taking the estimated aggregate shock distribution in S_t as given. Substituting for d_{it} from (5), the second expression inside the square brackets can be written as

$$\begin{aligned}
& \max_{\hat{k}_{it+1},\; b_{it+1}, w_{it+1}, l_{it+1}} \Pi_{it} - B_{it} - \hat{k}_{it+1} + \chi q_{it} b_{it+1} + \mathbb{E} M_{t+1} V\left(\Pi_{it+1}, B_{it+1}, S_{t+1}\right) \\
= \ & \Pi_{it} - B_{it} + \max_{\hat{k}_{t+1},\; b_{t+1}, w_{t+1}, l_{t+1}} -\hat{k}_{it+1} + \chi q_{it} b_{it+1} + \mathbb{E}_t M_{t+1} V\left(\Pi_{it+1}, B_{it+1}, S_{t+1}\right) \\
\equiv \ & \Pi_{it} - B_{it} + \Gamma_t
\end{aligned}$$

Finally, the value of the assets of a defaulting firm $\tilde{V}\left(\Pi_{it}, S_t\right)$ is simply the value of a firm with no external obligations, i.e. $V\left(\Pi_{it}, 0, S_t\right) = \tilde{V}\left(\Pi_{it}, S_t\right)$.

The two aggregate objects that are relevant for the firm's problem are the wage rate $\mathcal{W}_t$ and the stochastic discount factor, M_{t+1}. These depend on aggregate consumption and labor, implying that the aggregate state S_t can be summarized by $(\Pi_t, L_t, \hat{G}_t)$ where $\Pi_t \equiv AK_t^{\alpha}L_t^{1-\alpha} + (1-\delta)K_t$ is the aggregate resources availa.

For a given $\hat{G}_t$, an anticipated utility recursive equilibrium is a set of (i) functions for aggregate consumption and labor that maximize (1) subject to a budget constraint, (ii) firm value functions and associated policy functions that solve (7) , taking the bond price and wage functions (3) , and (6) and the stochastic discount factor (4) as given. (iii) aggregate consumption and labor are consistent with individual choices.

3 Solving the model

We now characterize the equilibrium and explore how tail events and the subsequent changes in beliefs affect the persistence and level of macro and financial outcomes. We present only the key equations here and relegate the detailed derivations to Appendix A.

We first note that, for any given aggregate capital shock ϕ_t, we can represent the optimal default policy as a threshold rule in the idiosyncratic output shock v_{it},

$$r_{it} = \begin{cases} 0 & \text{if} \quad v_{it} < \underline{v}(S_t) \\ 1 & \text{if} \quad v_{it} \geq \underline{v}(S_t) \end{cases}$$

Working from the first-order condition for the firm's capital choice, we find that the optimal $\hat{k}_{t+1}$ choice solves[9]

$$1 + \chi \mathcal{W}_t \frac{l_{t+1}}{\hat{k}_{t+1}} = \mathbb{E}\left[M_{t+1} R_{t+1}^k J^k(\underline{v})\right] \tag{8}$$

$$\begin{aligned} \text{where} \qquad R_{t+1}^k &= A\phi_{t+1}^{\alpha}\left(\frac{\hat{k}_{t+1}}{l_{t+1}}\right)^{\alpha-1} + (1-\delta)\,\phi_{t+1} \\ J^k(\underline{v}) &= 1 + h(\underline{v})(\theta\chi - 1) + \underline{v}(1 - F(\underline{v}))(\chi - 1) \\ h(\underline{v}) &\equiv \int_{-\infty}^{\underline{v}} v f(v) dv \end{aligned}$$

The term R_{t+1}^k is related to the return on capital, augmented with the capital quality shock, ϕ_{t+1}. The term $J^k(\underline{v})$ reflects the net effect of distortions induced by the tax advantage and default penalties associated with debt. In the absence of any these distortions (e.g. if $\chi = 1$), $J^k(\underline{v}) = 1$, reducing (8) to a standard Euler equation. In general, however, the wedge $J^k(\underline{v})$ distorts the equilibrium choice of capital away from the choices of a planner.

The optimality condition for labor looks quite similar. Just like with capital, firms equate

[9] Since all firms are identical, they make symmetric choices and accordingly, we suppress the i subscript.

the marginal cost of an additional unit of labor, namely $\mathcal{W}_t$, with the expected marginal product of labor, adjusted for the effect of additional promised wages on the cost of default:

$$\chi \mathcal{W}_t = \mathbb{E}\left[M_{t+1}\left(1-\alpha\right)A\phi_{t+1}^{\alpha}\left(\frac{\hat{k}_{t+1}}{l_{t+1}}\right)^{\alpha} J^{l}(\underline{v})\right] \tag{9}$$

$$\text{where} \qquad J^{l}(\underline{v}) = 1 + h\left(\underline{v}\right)\left(\theta\chi - 1\right) - \underline{v}^{2} f\left(\underline{v}\right)\chi\left(\theta - 1\right)$$

Finally, the firm's optimality condition with respect to leverage

$$\left(1-\theta\right)\mathbb{E}_t\left[M_{t+1}\underline{v}f\left(\underline{v}\right)\right] = \left(\frac{\chi-1}{\chi}\right)\mathbb{E}_t\left[M_{t+1}\left(1-F\left(\underline{v}\right)\right)\right] \tag{10}$$

The left hand side is the marginal cost of increasing leverage - it raises the expected losses from the default penalty (a fraction $(1-\theta)$ of the firm's value). The right hand side is the marginal benefit - the tax advantage times the value of debt issued.

The three optimality conditions, $(8)-(10)$, along with those from the household side - in particular, the labor supply condition (6) - characterize the equilibrium of this economy and can be solved numerically.

For numerical tractability, we make a simplifying assumption - instead of letting firms choose bond obligations period-by-period, we assume that they follow a simple rule and target a constant leverage (defined as the ratio of total obligations to capital). This is equivalent to replacing (10) with

$$\frac{B_{it+1}}{\hat{k}_{it+1}} = lev^{\text{Target}}$$

4 Measurement and Calibration

One of the key strengths of our belief-driven theory is that, by assuming that agents form beliefs as an econometrician would, we allow the data to discipline beliefs. In this section, we parameterize the model to match key features of the US economy. We then subject the model economy to the realized time series of capital quality shocks from US post-war data and evaluate the predictions for aggregates that we did not calibrate to, such as investment, output and consumption.

4.1 Measuring capital quality shocks

The next step is to construct a time series of $\{\phi_t\}$. We use annual data on non-financial assets of non-financial corporations in the US economy. The Flow of Funds reports published by the Federal Reserve contain two such series - one evaluated at historical cost and the other at replacement cost or market value. We interpret the latter as corresponding to effective capital. We then adjust for changes in the nominal price of investment goods, using the change in price index for non-residential investment from the National Income and Product Accounts [10], which allows us to recover the change in effective capital from the quality shock ϕ_t.

Formally, let

$$\begin{aligned} NFA_t^{RC} &= \text{Replacement cost of non-financial assets} \\ NFA_t^{HC} &= \text{Historical cost of non-financial assets} \\ PINDX_t^k &= \text{Investment price index (BEA)} \end{aligned}$$

We can map these objects into their model counterparts as follows:

$$\begin{aligned} P_t^k K_t &= NFA_t^{RC} \\ P_{t-1}^k \hat{K}_t &= (1-\delta) NFA_{t-1}^{RC} + P_{t-1}^k X_{t-1} \\ &= (1-\delta) NFA_{t-1}^{RC} + NFA_t^{HC} - (1-\delta)\, NFA_{t-1}^{HC} \end{aligned}$$

where X_{t-1} denotes investment in period $t-1$ and P_t the nominal price of capital goods. Then,

$$\phi_t = \frac{K_t}{\hat{K}_t} = \left(\frac{P_t^k K_t}{P_{t-1}^k \hat{K}_t}\right)\left(\frac{P_{t-1}^k}{P_t^k}\right) \tag{11}$$

$$= \left[\frac{NFA_t^{RC}}{(1-\delta) NFA_{t-1}^{RC} + NFA_t^{HC} - (1-\delta)\, NFA_{t-1}^{HC}}\right]\left(\frac{PINDX_{t-1}^k}{PINDX_t^k}\right) \tag{12}$$

where the second line uses $\left(\frac{PINDX_{t-1}^k}{PINDX_t^k}\right)$ as a proxy for $\left(\frac{P_{t-1}^k}{P_t^k}\right)$.

Using the measurement equation (12), we construct an annual time series for capital quality shocks for the US economy over the last few decades. The left panel of Figure 3 plots the resulting series. For most of the sample period, the shock realizations are in a relatively tight range around 1, but at the onset of the recent Great Recession, we saw two large adverse realizations: 0.93 in 2008 and 0.84 in 2009. To put this in context, the standard deviation of the series from 1950-2007 was 0.03. We will use the model to simulate the responses of the economy over time to negative shocks of this magnitude.

We then apply a kernel density estimation procedure to this time series to construct a sequence of beliefs. In other words, for each t, we construct $\{\hat{g}_t\}$ using the available time series until that point. The resulting estimates for two dates - 2007 and 2009 - are shown in the right

[10] Our results are robust to alternative measures of price changes.

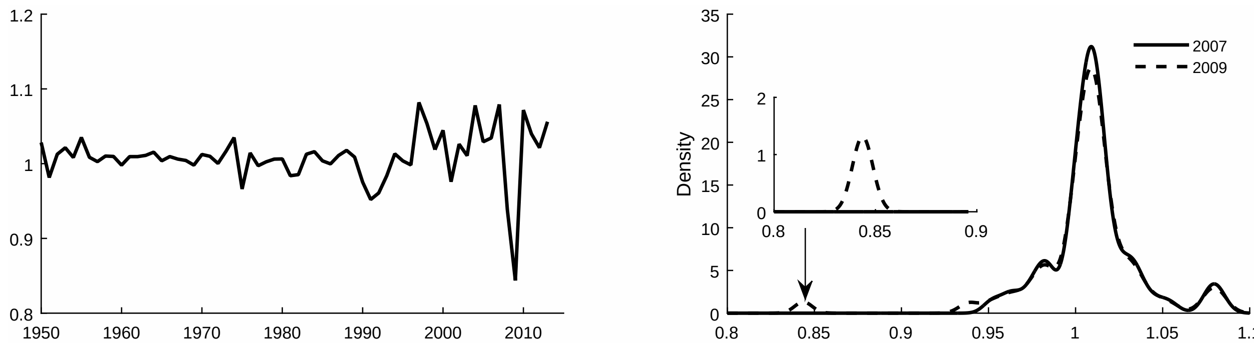

Figure 3: **Distribution of capital quality shocks.** The left tail shows the effect of the Great Recession.

panel of Figure 3. They show that the great recession induced a significant increase in the perceived tail risk. The density function for 2007 implies almost zero mass below 0.95, while the one for 2009 attach a non-trivial probability to significantly worse realizations.

4.2 Calibration

A period is interpreted as a year and the discount factor β is set to 0.9. The share of capital in the production, α, is set to 0.40. The recovery rate upon default, θ, is set to 0.70, following Gourio (2013). The distribution for the idiosyncratic shocks, v_{it} is assumed to be lognormal, i.e. $\ln v_{it} \sim N\left(-\frac{\hat{\sigma}^2}{2}, \hat{\sigma}^2\right)$ with $\hat{\sigma}^2$ chosen to target a default rate[11] of 0.02. The labor supply parameter, γ, is set to 0.5, in line with Midrigan and Philippon (2011), corresponding to a Frisch elasticity of 2. The labor disutility parameter ζ and the TFP term in production are normalized to 1.

For the parameters governing risk aversion and intertemporal elasticity of substitution, we use standard values from the asset pricing literature[12] and set $\psi = 0.5$ (or equivalently, an IES of 2) and $\eta = 10$. The leverage target is 0.70, obtained by adding the wage bill (approximately 0.2 of the steady state capital stock) to financial leverage (the ratio of external debt to capital, which stands at roughly 0.5 in US data - from Gourio (2013)). Since leverage is exogenous, the tax advantage χ is a free parameter. We set it to a baseline value of 1.06 and verified numerically that our results are not particularly sensitive to this choice.

Table 1 summarizes all parameter choices.

[11]This is in line with the target in Khan et al. (2014), though a bit higher than the one in Gourio (2013). We verified that our quantitative results are not sensitive to this target.

[12]See discussion in Gourio (2013).

Parameter	Value	Description
β	0.91	Discount factor
η	10	Risk aversion
ψ	0.50	1/Intertemporal elasticity of substitution
γ	0.50	1/Frisch elasticity
ζ	1	Labor disutility
α	0.40	Capital share
δ	0.03	Depreciation rate
A	1	TFP
χ	1.06	Tax advantage of debt
θ	0.70	Recovery rate
$\hat{\sigma}$	0.33	Idiosyncratic volatility
lev^{Target}	0.70	Leverage ratio

Table 1: **Parameters**

5 Quantitative Results

Our main goal in this section is to quantify the size and persistence of the macroeconomic response to a large but transitory shock ϕ_t and explore the role of belief updating, debt and risk aversion to this response. With this goal in mind, we perform the following experiment using historical data on ϕ_t realizations from 1950-2009, measured using the strategy outlined in Section 4. We begin by estimating $\hat{G}_{2007}$ using the data through 2007. Then, starting from the steady state associated with this estimated distribution[13], we subject the model economy to two adverse realizations - 0.93 and 0.84, which correspond to the shocks that we observed in 2008 and 2009. This leads to a revised estimate for the distribution, $\hat{G}_{2009}$. At this point, we could keep adding 2010-14 data, re-estimating each period (we do this in the Appendix). We could even simulate data beyond 2014 by drawing from the estimated distribution. But we know that beliefs are martingales. The average distribution estimated on draws from $\hat{G}_{2009}$ is $\hat{G}_{2009}$. Therefore, we hold beliefs to be $\hat{G}_{2009}$ and assume that the shock realizations from 2010 on are equal to their average value. This exercise isolates the persistent response to an isolated shock to beliefs, without mixing up that endogenous response with the effect of additional exogenous shocks.[14]

The resulting impulse response functions are shown in Figure 4. The top left panel shows the

[13] The steady state is obtained by simulating the model for 1000 periods using the $\hat{G}_{2007}$ and the associated policy functions, discarding the first 500 observations and time-averaging across the remaining periods.

[14] The impulse responses from this experiment are a reasonable an approximation to those where the economy is subjected to draws from $\hat{G}_{2009}$ and beliefs are updated with each new draw. Impulse response functions can then be derived by averaging across sample paths. In the Appendix, we perform this exercise for the version of our model without debt (which is a lot more tractable) and show that it yields almost identical impulse response functions as the experiment holding beliefs fixed.

time path for ϕ_t (as deviations from its average value). The remaining panels show the behavior of output, capital and employment over time. They show that the negative realizations in 2008 and 2009 and the resulting belief revisions induce a prolonged stagnation, with the economy trending towards a new, lower (stochastic) steady state. Output in this new steady state is 13% lower than the associated one under $\hat{G}_{2007}$. The corresponding drops in capital and labor are about 19% and 9% respectively. Thus, even though the ϕ_t shocks were transitory, they were so large that the resulting change in beliefs permanently reduces economic activity

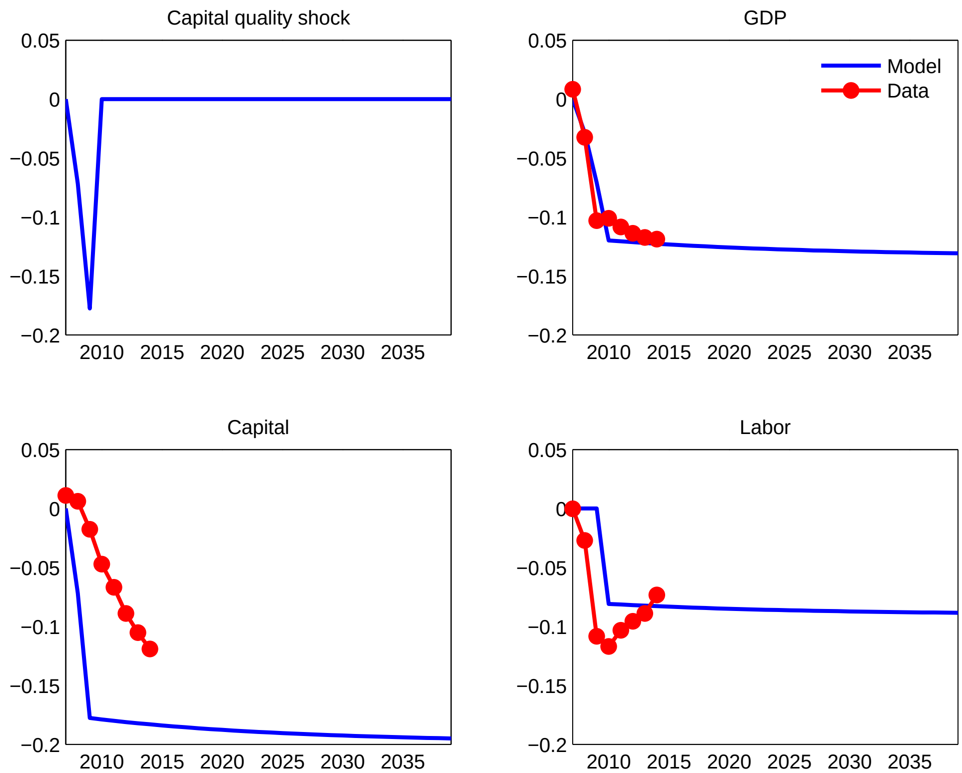

Figure 4: **Large negative shocks create extremely persistent responses in output, investment and labor.** Solid line shows the change in aggregates (relative to the stochastic steady state associated with $\hat{G}_{2007}$). The circles show de-trended US data for the period 2009-2014.

Figure 4 also plots the actual data on the output, capital stock and labor (in deviations from their respective pre-2007 trends) for the US economy.[15] As the graph shows, the model's predictions for the drop in output line up remarkably well with the data. The predicted path for capital and labor are also similar to the observed patterns though there are some differences. The former exhibits a sharper drop than the observed drop in capital input - which is not particularly surprising, given that our model abstracted from adjustment costs and other frictions that could induce a more sluggish response of capital. Similarly, the model's predictions for

[15] We use data on output, capital and labor input from Fernald (2014). Each series is adjusted for growth in working age population and then detrended using a log-linear trend estimated using data from 1950-2007.

	2015	Long run
Data	-0.13	–
Benchmark model	-0.13	-0.13
Model without learning	-0.09	0.00
Model without debt	-0.13	-0.06
Model with quasi-linear preferences	-0.07	-0.07

Table 2: Change in GDP (relative to 2007 steady state).

labor underpredict the actual change in employment. In the data, employment dropped sharply in 2008-'09, almost contemporaneously with the negative shocks and then recovered slowly. In the model, however, the drop occurs later, but that is largely due to the assumption that labor is chosen in advance. Bringing the model closer to the data along these dimensions is no doubt important and will require a richer model with additional features and frictions, but Figure 4 demonstrates the quantitative potential of learning in a standard business cycle setting[16].

5.1 Discussion

We now perform a series of counterfactual experiments with our calibrated model to provide a deeper understanding of the effects of tail events on economic outcomes. These exercises will allow us to isolate the various forces that combine to deliver the results presented in Figure 4.

The first experiment focuses on the changes in beliefs by comparing our baseline results to an economy with no learning. It also sheds light on the differential impact of large vs small shocks and role of higher moments. The second experiment analyzes how learning interacts with debt by comparing our results to an economy with no leverage. Finally, we quantify the contribution of risk aversion by analyzing a version of our model with quasi-linear preferences. Table 2 summarizes our main conclusions. It shows that all 3 components - learning, debt and risk aversion play a significant role in generating persistent stagnation.

5.1.1 Role of Learning

Turning off belief revisions: In Figure 5, we compare our results to an otherwise identical economy where agents are assumed to know the final distribution $\hat{G}_{2009}$ throughout from the very beginning and so, do not revise their beliefs. This corresponds to a standard rational expectations approach, where agents are assumed to know the true distribution. The econometrician estimates this distribution using all the available data. The solid and dashed lines

[16] In the Appendix, we show that including the effect of shock realizations post-2009 does not materially change this finding.

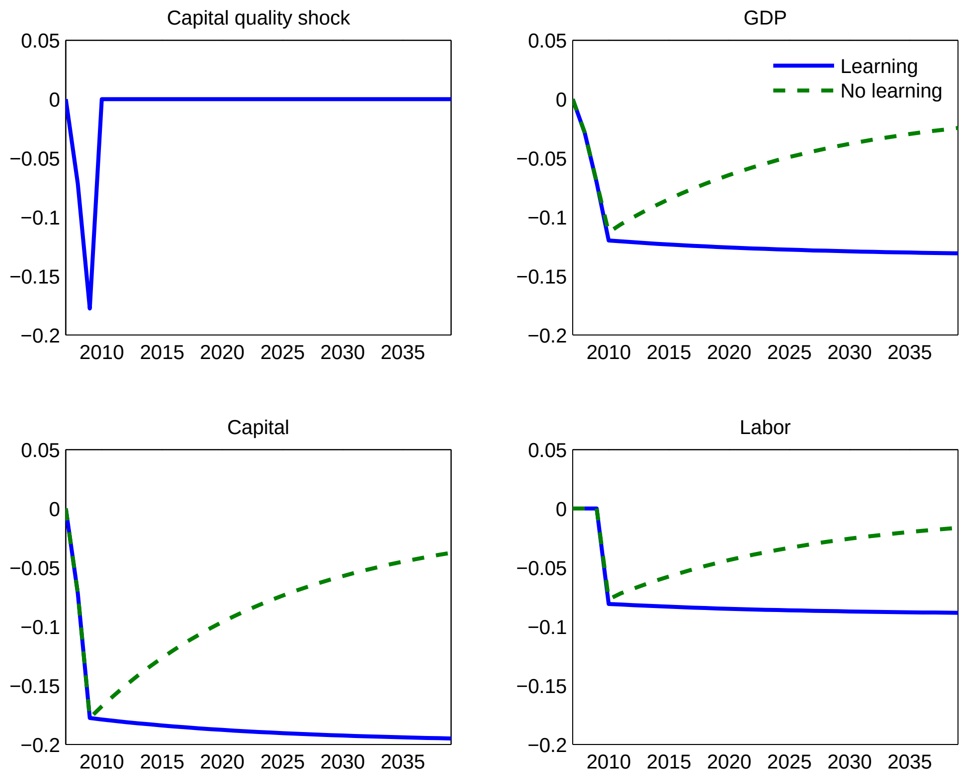

Figure 5: **Fixing beliefs creates sharp rebounds in output, investment and labor.** Solid line (learning) is the model with belief updating. Dashed line (no learning) is an identical model where agents believe that shocks are drawn from the distribution estimated on the full sample of data and never revise those beliefs. Zero is the steady state level in each economy.

in the remaining panels show the response of aggregate variables with and without learning respectively. The impulse-response functions are computed as before. Each economy is assumed to be at its stochastic steady state in 2007 and is subjected to the same sequence of shocks - two large negative ones in 2008 and 2009 and the average level subsequently.

They show that, in the absence of belief revisions, the negative shock prompts firms to increase investment to replenish the lost effective capital. While the curvature in the utility function moderates the speed of this transition to an extent, the overall pattern of a steady recovery back to the original steady state is clear[17]. This shows that learning is central to the model's ability to generate long-lived reductions in economic activity.

Large vs small shocks: Recall from Figure 3 that the adverse shocks observed during the 2008-'09 period were exceptionally large - the shock in 2009, for example, was almost 5

[17] Since the no-learning economy is endowed with the same end-of-sample beliefs as the learning model, they both ultimately converge to the same *levels*. But they start at different steady states (normalized to 0 for each series).

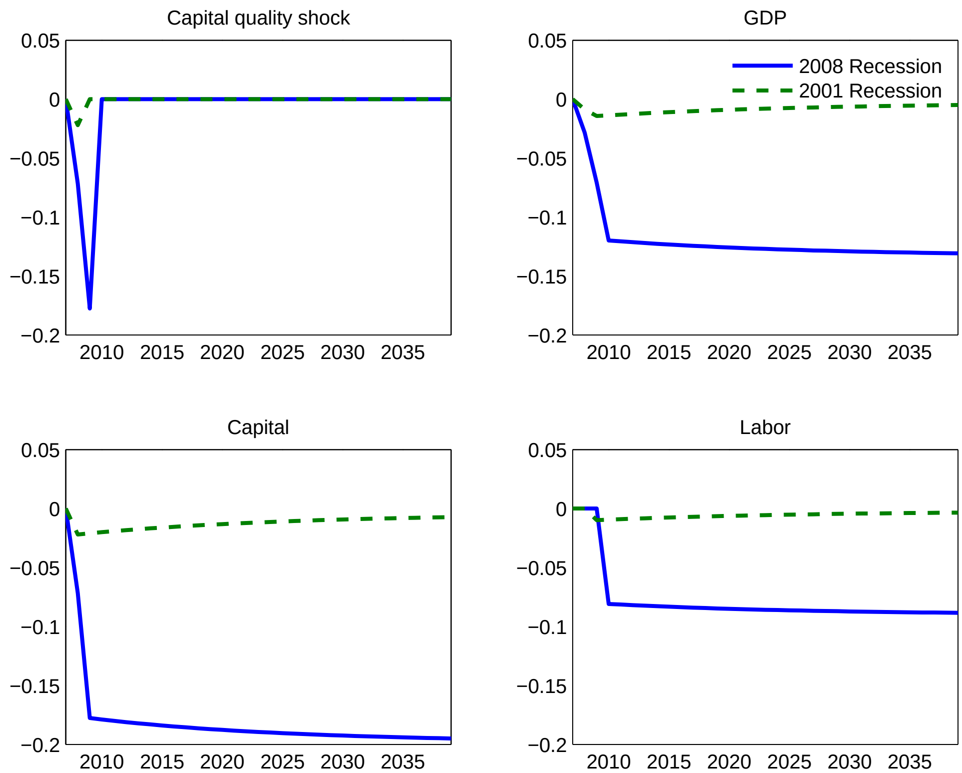

Figure 6: **Small shocks generate transitory effects.** Solid line (2008 recession) is the model with shocks equal to those observed during 2008-09. Dashed line (2001 recession) is the counter-factual simulation with a 1 σ shock.

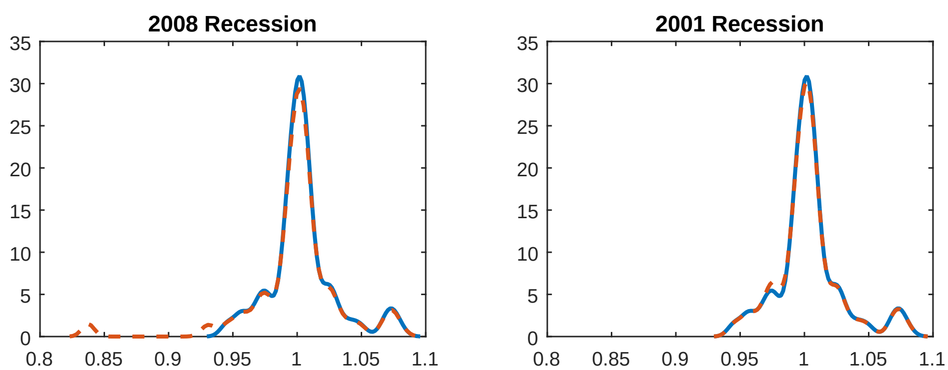

Figure 7: **Small shocks imply negligible belief revisions.** The solid blue line in both panels shows the estimated density in 2007, while the dashed lines show the new estimate after the 2008-09 shocks (left panel) and a counter-factual 1σ shock (the right panel).

standard deviations below the mean. To better understand the effect of such a large shock, we conduct a counter-factual simulation and subject the model to a much smaller shock - 1 standard deviation below the mean, which is roughly in line with what we observed during the 2001-'02 recession. The results are shown in Figure 6. The 1σ shock has a much smaller effect on impact, but perhaps more interestingly, the effects of aggregate outcomes are transitory, unlike those arising from the 2008-'09 shock. Both shocks induce permanent belief revisions but the magnitude of these changes are much smaller under the 1σ shock. As a result, the new stochastic steady state is not that much different from the starting point, causing the economy to return, albeit slowly, to more or less the same level of economic activity as before the shock. With the 2008-09 shock, however, the change in beliefs (and through them, on aggregates) is quite dramatic, leading to very different long run outcomes. This experiment reveals the ability of our learning model to rationalize why some recessions (like the Great Recession) can look very different from earlier episodes[18].

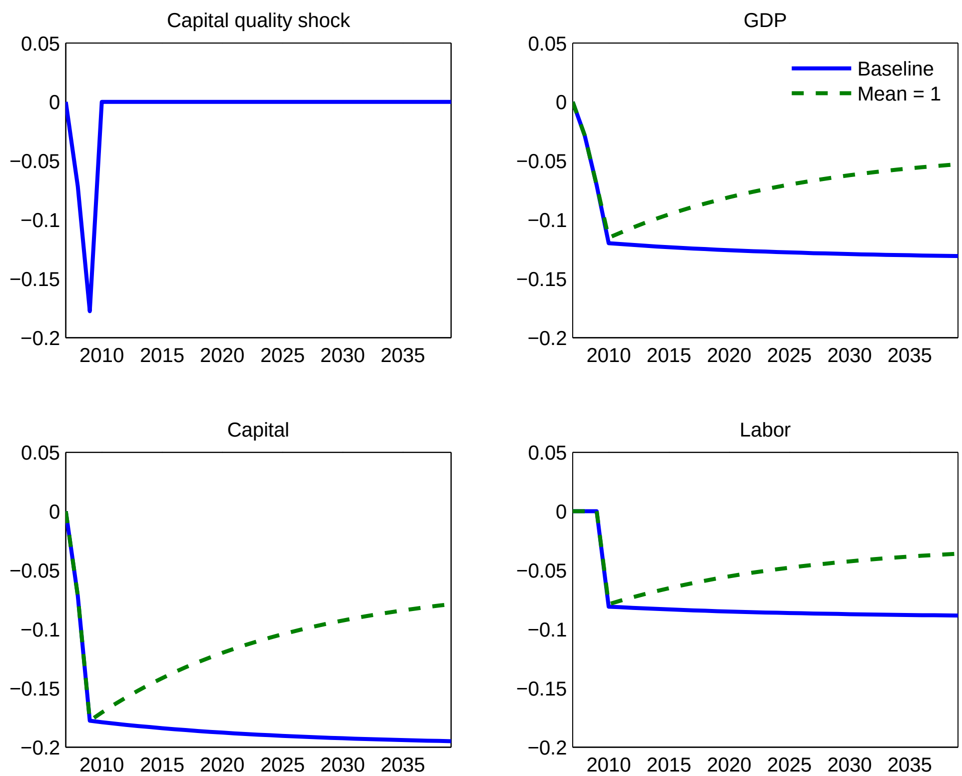

Figure 8: **Higher moments account for almost half of the long-run effects.** Solid line (Baseline) is the model in section 2, where beliefs are updated according to (2). Dashed line (Mean = 1) is an identical model where the estimation distribution $\hat{G}_t$ is re-scaled so that $\mathbb{E}_t(\phi_t) = 1$. Zero is the initial steady state level in each economy.

Mean beliefs vs higher moments: As we saw from Figure 3, the Great Recession had a pronounced impact on the perceived distribution of shocks. Here, we decompose the total

[18]Note that this is not simply a statement about large vs small shocks - what matters is not the size of the shock *per se* but the effect it has on beliefs. For example, large shocks also may have transitory effects in an economy where such shocks have been observed very frequently.

effect into a component attributable to changes in the mean and the remaining attributable to changes in higher moments. To do this, we re-scale the estimated distributions at each t so that $\mathbb{E}_t(\phi_t) = 1$. This isolate the effect of higher moments (as opposed to the mean). The dashed lines in Figure 8 show the impulse response functions under this modification. They show that changes in the mean and higher moments are equally important in generating the persistent decline in economic activity. The long-run drop in GDP attributable to higher moments, for example, is about 6%, about half of the total effect.

5.1.2 Belief revisions and debt

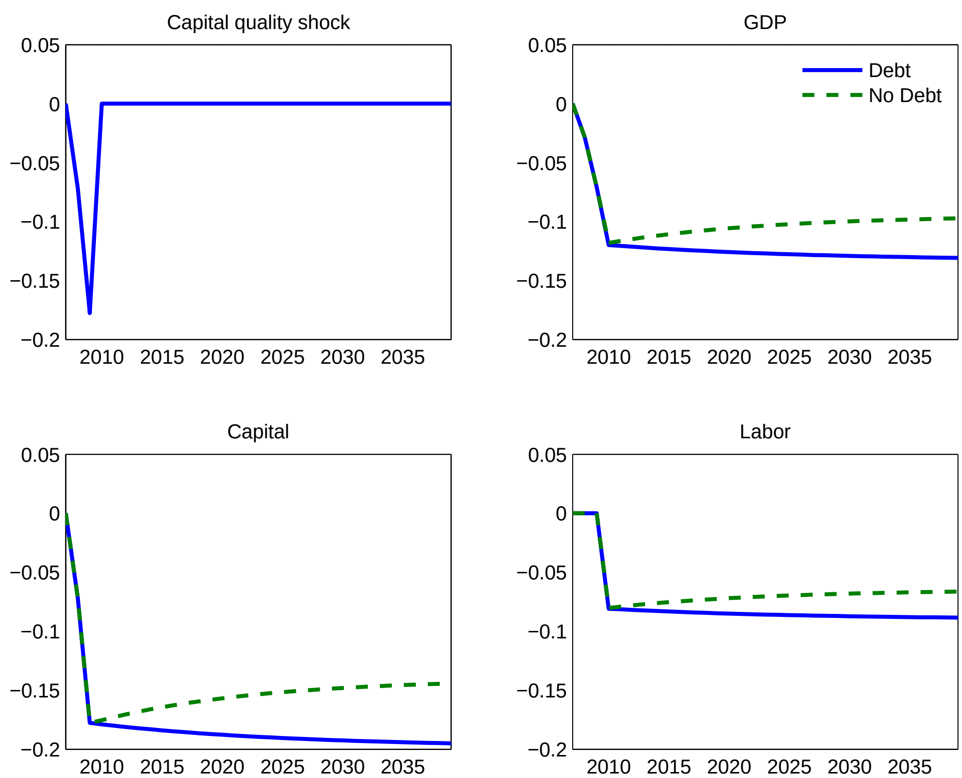

Figure 9: **Debt amplifies the effects of belief revisions on output, investment and labor.** Solid line (debt) is the model in section 2. Dashed line (no debt) is an identical model with $\chi = 0$, where firms choose zero debt. Zero is the steady state level in each economy.

Next, we examine the interaction between learning and debt by comparing our results to an identical economy where all investment is financed through equity (beliefs are updated over time, exactly as in our baseline model). Formally, we set the tax advantage parameter χ to 1 and the leverage target to 0. This implies that $J^k(\underline{v}) = J^l(\underline{v}) = 1$, i.e. the debt-related distortions in capital and labor choice disappear.

Figure 9 plots the time path for aggregate variables for this variant of our model, along with our baseline version from Figure 4. The graph shows that the effects of belief revisions are smaller in the absence of debt - for example, they lead to a 9% reduction in output (compared to 13% in the baseline version with debt). Thus, defaultable debt amplifies the effects of changes in tail risk and contributes almost a third of the long run macroeconomic response.

5.1.3 Belief revisions and risk aversion

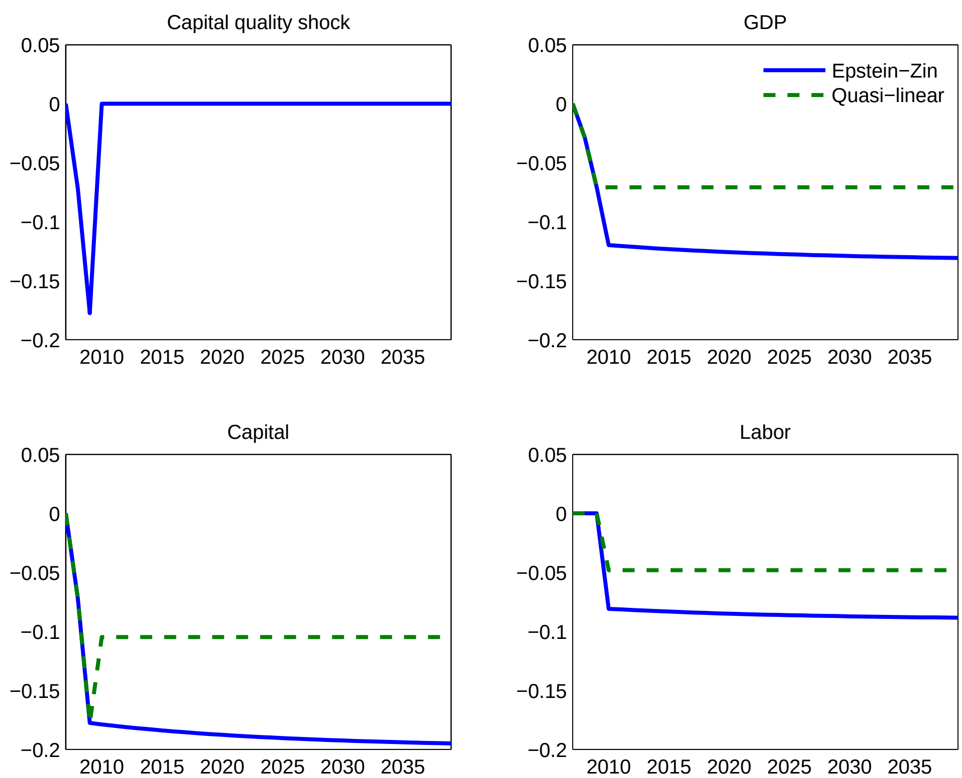

Figure 10: **Without curvature, transitions are starker and the drop in long-run GDP is half of our baseline effects.** Solid line (Epstein-Zin) is the baseline calibration with EZ preferences. Dashed line (Quasi-linear) is an identical model with $\psi = \eta = 0$.

Finally, we investigate the implications of risk aversion by comparing our results to an otherwise identical economy with quasilinear preferences. Formally, we set $\psi = \eta = 0$, so the utility function of the representative household reduces to $C_t - \zeta \frac{L_t^{1+\gamma}}{1+\gamma}$. This eliminates the desire for consumption smoothing and risk premia. While these elements are no doubt important in reality, suppressing them allows us to see how much of the persistent drop in investment and hiring comes from changes in the perceived distribution of returns to investing.

Figure 10 presents the time path for aggregate variables in this version, both with and without learning. As we would expect, the absence of curvature in consumption means that the economy transitions immediately to the new steady state. However, belief revisions still have substantial, permanent effects on the level of economic activity. For example, they lead to a drop in steady state output of about 7%. With risk aversion, this drop is almost doubled. This is because now capital (and labor) have to earn a risk premium, which changes with beliefs, and is particularly sensitive to tail risk. This further dampens firms' incentives to invest (and hire) after observing a large negative shock. The graph indicates that this effect is quite strong and accounts for almost half of the long-run drop in output.

6 Conclusion

No one knows the true distribution of shocks to the economy. Economists typically assume that agents in their models do know this distribution as a way to discipline beliefs. But assuming that agents do the same kind of real-time estimation that an econometrician would do is equally disciplined and more plausible. For many applications, assuming full knowledge has little effect on outcomes and offers tractability. But for outcomes that are sensitive to tail probabilities, the difference between knowing these probabilities and estimating them with real-time data can be large. The estimation error can be volatile and can introduce new, persistent dynamics into a model with otherwise transitory shocks. The essence of the persistence mechanism is this: Once observed, a shock (a piece of data) stays in one's data set forever and therefore permanently affects belief formation.

When firms finance investments with debt, they make investment and output sensitive to tail risk. Debt is an asset whose payoffs are flat throughout most of the state space, but very sensitive to the state for left-tail, default events. Therefore, the cost of debt depends precisely on the probabilities of a tail event, which are hardest to estimate and whose estimates fluctuate greatly. When debt (leverage) is low, the economy is not very sensitive to tail risk, and economic shocks will be more transitory. The combination of high debt levels and a shock that is a negative outlier makes tail risk surge, investment fall and depresses output in a persistent way.

When we quantify this mechanism and use capital price and quantity data to directly estimate beliefs, our model's predictions resemble observed macro outcomes in the wake of the great recession. These results suggests that perhaps persistent stagnation arose because, after seeing how fragile our financial sector is, market participants will never think about tail risk in the same way again.

References

Angeletos, G.-M. and J. La'O (2013). Sentiments. *Econometrica 81*(2), 739–779.

Backus, D., A. Ferriere, and S. Zin (2015). Risk and ambiguity in models of the business cycle. *Journal of Monetary Economics forthcoming.*

Beaudry, P. and F. Portier (2004). An exploration into pigou's theory of cycles. *Journal of Monetary Economics 51*, 1183–1216.

Bloom, N., M. Floetotto, N. Jaimovich, I. Sapora-Eksten, and S. Terry (2014). Really uncertain business cycles. NBER working paper 13385.

Brunnermeier, M. and Y. Sannikov (2014). A macroeconomic model with a financial sector. *American Economic Review 104*(2).

Cogley, T. and T. Sargent (2005). The conquest of us inflation: Learning and robustness to model uncertainty. *Review of Economic Dynamics 8*, 528–563.

Eaton, J. and M. Gersovitz (1981). Debt with potential repudiation: Theoretical and empirical analysis. *The Review of Economic Studies*, 289–309.

Eggertsson, G. and N. Mehrotra (2014). A model of secular stagnation. NBER Working Paper No. 20574.

Fajgelbaum, P., E. Schaal, and M. Taschereau-Dumouchel (2014). Uncertainty traps. NBER Working Paper No. 19973.

Fernald, J. (2014). A quarterly, utilization-adjusted series on total factor productivity. *Federal reserve bank of San Francisco working paper 19*.

Gertler, M. and P. Karadi (2011). A model of unconventional montary policy. *Journal of Monetary Economics 58*, 17–34.

Gertler, M., N. Kiyotaki, et al. (2010). Financial intermediation and credit policy in business cycle analysis. *Handbook of monetary economics 3*(3), 547–599.

Gourio, F. (2012). Disaster risk and business cycles. *American Economic Review 102(6)*, 2734–66.

Gourio, F. (2013). Credit risk and disaster risk. *American Economic Journal: Macroeconomics 5(3)*, 1–34.

Hall, R. (2014). Quantifying the lasting harm to the u.s. economy from the financial crisis. *NBER Macroeconomics Annual 29*, forthcoming.

Hansen, L. (2007). Beliefs, doubts and learning: Valuing macroeconomic risk. *American Economic Review 97(2)*, 1–30.

Huo, Z. and N. Takayama (2015). Higher order beliefs, confidence, and business cycles. University of Minnesota working paper.

Jaimovich, N. and S. Rebelo (2006). Can news about the future drive the business cycle? *American Economic Review 99(4)*, 1097–1118.

Johannes, M., L. Lochstoer, and P. Collin-Dufresne (2015). The asset pricing implications of priced structural parameter uncertainty.

Johannes, M., L. Lochstoer, and Y. Mou (2015). Learning about consumption dynamics. *Journal of Finance forthcoming.*

Khan, A., T. Senga, and J. K. Thomas (2014). Default risk and aggregate fluctuations in an economy with production heterogeneity. Technical report, Ohio State University working paper.

Kreps, D. (1998). *Anticipated Utility and Dynamic Choice*. Cambridge University Press.

Lorenzoni, G. (2009). A theory of demand shocks. *American Economic Review 99*(5), 2050–84.

Midrigan, V. and T. Philippon (2011). Household leverage and the recession.

Moriera, A. and A. Savov (2015). The macroeconomics of shadow banking. New York University working paper.

Navarro, G. (2014). Financial crises and endogenous volatility. *manuscript, New York University*.

Nimark, K. (2014). Man-bites-dog business cycles. *American Economic Review 104(8)*, 2320–67.

Orlik, A. and L. Veldkamp (2014). Understanding uncertainty shocks and the role of the black swan. Working paper.

Piazzesi, M., J. Salomao, and M. Schneider (2015). Trend and cycle in bond risk premia. Stanford University working paper.

Pintus, P. and J. Suda (2015). Learning financial shocks and the great recession. Aix Marseille working paper.

Straub, L. and R. Ulbricht (2013). Credit crunches, information failures, and the persistence of pessimism. Toulouse School of Economics working paper.

Van Nieuwerburgh, S. and L. Veldkamp (2006). Learning asymmetries in real business cycles. *Journal of Monetary Economics 53(4)*, 753–772.

Veldkamp, L. and J. Wolfers (2007). Aggregate shocks or aggregate information? costly information and business cycle comovement. *Journal of Monetary Economics 54*, 37–55.

Appendix

A Optimality conditions from firm's problem

The firm's optimization problem is

$$\begin{aligned}
V\left(k_{it}, l_{it}, w_{it}, b_{it}, S_t\right) &= \max\left[0,\ \ \Pi_{it} - B_{it} + \Gamma_{it}\right] \\
\Gamma_{it} &= \max_{\hat{k}_{it+1},\ b_{it+1}, w_{it+1}, l_{it+1}} -\hat{k}_{it+1} + \chi q b_{it+1} + \mathbb{E} M_{t+1} V\left(k_{it+1}, l_{it+1}, w_{it+1}, b_{it+1}, S_{t+1}\right) \\
\Pi_{it} &= v_{it}\left(A(\phi_t k_{it})^{\alpha} l_{it}^{1-\alpha} + (1-\delta)\phi_t k_{it}\right) \\
B_{it+1} &= b_{it+1}\ + w_{it+1} l_{it+1} \\
q\left(\hat{k}_{it+1},\ l_{it+1}, B_{it+1}, S_t\right) &= \mathbb{E}_t M_{t+1}\left[r_{it+1} + (1-r_{it+1})\frac{\theta V\left(k_{it}, l_{it}, 0, 0, S_t\right)}{B_{it+1}}\right] \\
w_{it+1} q &= \mathcal{W}_t \\
r_{it+1} &= \begin{cases} 0 & \text{if} \quad v_{it} < \underline{v}\left(S_t\right) \\ 1 & \text{if} \quad v_{it} \geq \underline{v}\left(S_t\right) \end{cases} .
\end{aligned}$$

First, note that we can write the firm's problem in term of leverage and labor capital ratio, defined as $lev_{it+1} \equiv \frac{B_{it+1}}{\hat{k}_{it+1}}$ and $\frac{l_{it+1}}{\hat{k}_{it+1}}$. Then,

$$R^k\left(\frac{l_{it+1}}{\hat{k}_{it+1}}, \phi_{t+1}\right) \equiv \frac{\Pi_{it+1}}{\hat{k}_{it+1}} = v_{it}\left(A({\phi_{t+1}}^{\alpha}\left(\frac{l_{it+1}}{\hat{k}_{it+1}}\right)^{1-\alpha} + (1-\delta)\phi_{t+1}\right).$$

This implies that

$$\Gamma_{it} = \max_{\hat{k}_{it+1},\ lev_{it+1}, \frac{l_{it+1}}{\hat{k}_{it+1}}} \hat{k}_{it+1}\left(-1 - \chi\mathcal{W}_t\frac{l_{it+1}}{\hat{k}_{it+1}} + \chi q lev_{it+1}\ + \mathbb{E} M_{t+1} r_{t+1}\left(v_{it} R^k_{t+1} - lev_{it+1} + \frac{\Gamma_{it+1}}{\hat{k}_{it+1}}\right)\right)$$

$$q\left(\frac{l_{it+1}}{\hat{k}_{it+1}}, lev_{it+1}, S_t\right) = \mathbb{E}M_{t+1}\left[r_{t+1} + (1-r_{t+1})\,\theta\frac{v_{it}R^k_{t+1} + \frac{\Gamma_{it+1}}{\hat{k}_{it+1}}}{lev_{it+1}}\right].$$

We guess and then verify that $\Gamma_{it+1} = 0$.[19] Replacing the debt price schedule and rearranging terms yields

$$\Gamma_{it} = \max_{\hat{k}_{it+1},\ lev_{it+1}, \frac{l_{it+1}}{\hat{k}_{it+1}}} \hat{k}_{it+1}\left(-1 - \chi\mathcal{W}_t\frac{l_{it+1}}{\hat{k}_{it+1}} + \mathbb{E}M_{t+1}\tilde{J}^k_{t+1}\right)$$

$$\tilde{J}^k_{t+1} = R^k_{t+1} + lev_{it+1}(\chi-1)\,r_{t+1} + (\chi\theta-1)(1-r_{t+1})\,v_{it}R^k_{t+1}.$$

The expectation with respect to the idiosyncratic shock implies $\mathbb{E}r_{t+1} = (1-F(\underline{v}))$. Also, note that the default threshold becomes $\underline{v} = \frac{lev_{it+1}}{R^k_{t+1}}$. Hence

$$\tilde{J}^k_{t+1} = R^k_{t+1}\left(1 + \underline{v}(\chi-1)(1-F(\underline{v})) + (\chi\theta-1)\,h(\underline{v})\right)$$

where $h(v) = \int_{-\infty}^{v} vdF(v)$. Finally, the problem is

$$\Gamma_{it} = \max_{\hat{k}_{it+1},\ lev_{it+1}, \frac{l_{it+1}}{\hat{k}_{it+1}}} \hat{k}_{it+1}\left(-1 - \chi\mathcal{W}_t\frac{l_{it+1}}{\hat{k}_{it+1}} + \mathbb{E}M_{t+1}R^k_{t+1}J^k(\underline{v})\right)$$

$$J^k(\underline{v}) = 1 + (\chi-1)\,\underline{v}\,(1-F(\underline{v})) + (\chi\theta-1)\,h(\underline{v})$$

$$\underline{v} = \frac{lev_{it+1}}{R^k_{t+1}}$$

First, note that the problem is linear in $\hat{k}_{it+1}$ therefore in equilibrium we must have that

$$1 + \chi\mathcal{W}_t\frac{l_{it+1}}{\hat{k}_{it+1}} = \mathbb{E}M_{t+1}R^k_{t+1}J^k(\underline{v}),$$

which implies equation 8 in the main text and in turn it verifies the guess, $\Gamma_{it} = 0$.

Next, the first order condition with respect to $\frac{l_{t+1}}{\hat{k}_{it+1}}$ is

$$\chi\mathcal{W}_t = \mathbb{E}M_{t+1}R^k\frac{\partial J^k(\underline{v})}{\partial\frac{l_{t+1}}{\hat{k}_{it+1}}} + \mathbb{E}M_{t+1}\frac{\partial R^k}{\partial\frac{l_{t+1}}{\hat{k}_{it+1}}}J^k(\underline{v}),$$

[19] As the firm has constant returns to scale the problem will be linear in capital and in equilibrium $\Gamma_{it} = 0$. See Navarro (2014).

where

$$
\begin{aligned}
R^k_{t+1}\frac{\partial J^k(\underline{v})}{\partial \frac{l_{t+1}}{\hat{k}_{it+1}}} &= R^k_{t+1}\frac{\partial \underline{v}}{\partial \frac{l_{t+1}}{\hat{k}_{it+1}}}\left((\chi-1)\left(1-F\left(\underline{v}\right)\right)-\underline{v}\left(\chi-1\right)f\left(\underline{v}\right)+\left(\chi\theta-1\right)\frac{\partial h\left(\underline{v}\right)}{\partial \underline{v}}\right)\\
\frac{\partial \underline{v}}{\partial \frac{l_{t+1}}{\hat{k}_{it+1}}} &= -\frac{lev_{it+1}}{\left(R^k\right)^2}\frac{\partial R^k}{\partial \frac{l_{t+1}}{\hat{k}_{it+1}}} = -\frac{\underline{v}^2}{lev_{it+1}}\frac{\partial R^k}{\partial \frac{l_{t+1}}{\hat{k}_{it+1}}}\\
\frac{dh\left(\underline{v}\right)}{d\underline{v}} &= \underline{v}f\left(\underline{v}\right)\\
\frac{\partial R^k_{t+1}}{\partial \frac{l_{t+1}}{\hat{k}_{it+1}}} &= v_{it}A\left(1-\alpha\right)\phi_{t+1}{}^{\alpha}\left(\frac{l_{it+1}}{\hat{k}_{it+1}}\right)^{-\alpha}.
\end{aligned}
$$

Rearranging terms yields

$$
\begin{aligned}
\chi\mathcal{W}_t &= \mathbb{E}M_{t+1}\frac{\partial R^k}{\partial \frac{l_{t+1}}{\hat{k}_{it+1}}}J^l\left(\underline{v}\right)\\
J^l\left(\underline{v}\right) &= 1+\underline{v}^2 f\left(\underline{v}\right)\chi\left(1-\theta\right)-\left(1-\chi\theta\right)h\left(\underline{v}\right),
\end{aligned}
$$

which is (9) in the main text.

Finally, the first order condition with respect to leverage is

$$
\mathbb{E}M_{t+1}R^k_{t+1}\frac{\partial J^k_{t+1}}{\partial lev_{it+1}} = 0,
$$

where

$$
\begin{aligned}
\frac{\partial J^k_{t+1}}{\partial lev_{it+1}} &= \frac{\partial \underline{v}}{\partial lev_{it+1}}\left(\left(\chi-1\right)\left(1-F\left(\underline{v}\right)\right)-\left(\chi-1\right)\underline{v}f\left(\underline{v}\right)+\left(\chi\theta-1\right)\underline{v}f\left(\underline{v}\right)\right)\\
&= \frac{1}{R^k_{t+1}}\left(\left(\chi-1\right)\left(1-F\left(\underline{v}\right)\right)-\chi\left(1-\theta\right)\underline{v}f\left(\underline{v}\right)\right)
\end{aligned}
$$

hence

$$
\left(1-\theta\right)\mathbb{E}_t\left[M_{t+1}\underline{v}f\left(\underline{v}\right)\right] = \left(\frac{\chi-1}{\chi}\right)\mathbb{E}_t\left[M_{t+1}\left(1-F\left(\underline{v}\right)\right)\right],
$$

which is (10) in the main text.

B Belief revisions post-2009

Figure 4 in the text was generated by fixing post-2009 shock realizations (at the average) and beliefs (at $\hat{G}_{2009}$). In this section, we show that our results are not particularly sensitive to these assumptions, which were made purely for tractability. For this purpose, we use the version of the model without debt ($\chi = 1, lev^{Target} = 0$), which is computationally much more tractable. Specifically, for periods after 2009, we subject the economy to draws from $\hat{G}_{2009}$ and beliefs are updated with each new draw. Impulse response functions are derived by averaging across sample paths. Figure 11 shows the results, along with those that emerge from the fixed beliefs formulation of this version of the model. It shows that fixing beliefs approximates the full blown simulation remarkably well.

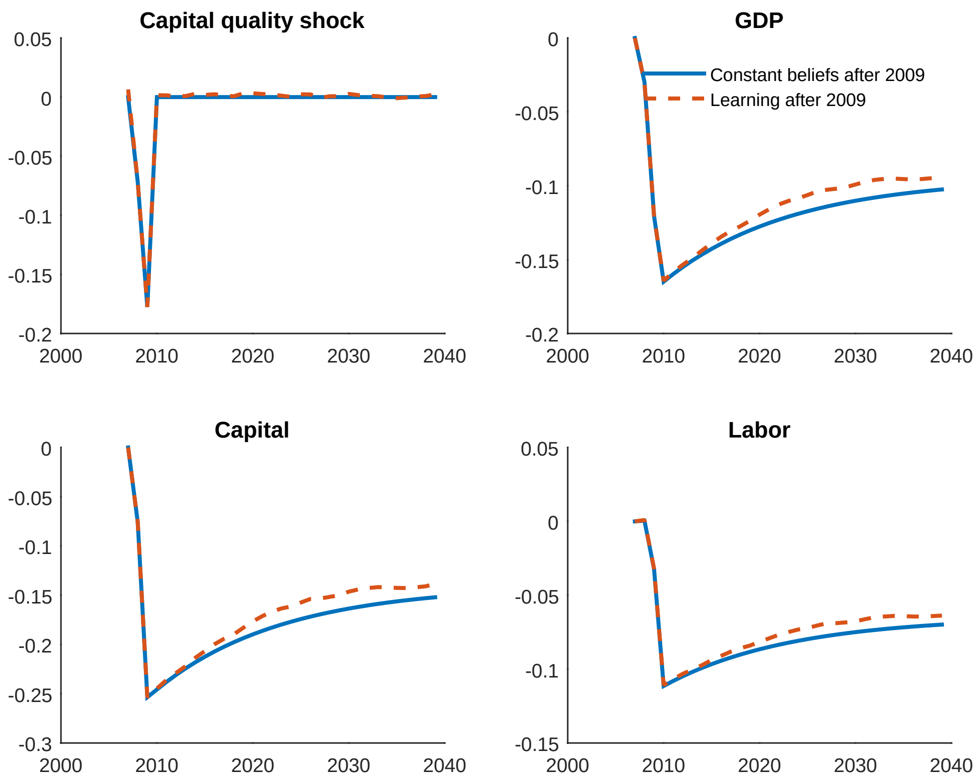

Figure 11: **Model vs data from 2008-2014.** Solid line is the baseline model subjected to the observed sequence of shocks from 2008-2014. The red circles are US data, in deviations from their pre-crisis trends.

C Effect of shocks 2010-2014

Here, we subject our baseline calibrated model to the full sequence of shocks, from 2008 through 2014. Agents' decisions in each year are a function of the appropriate estimated distribution, in line with our anticipated utility framework. For the period after 2014, we adopt the same strategy as in Figure 4, i.e. we hold fixed the shock realizations (at their average) and beliefs (at $\hat{G}_{2014}$). The resulting time paths are plotted in Figure 12, along with the de-trended data. The patterns implied by the model are quite close to the observed, particularly for output and employment. As before, the model overshoots in terms of the capital response, but that could be addressed by incorporating adjustment costs or other frictions.

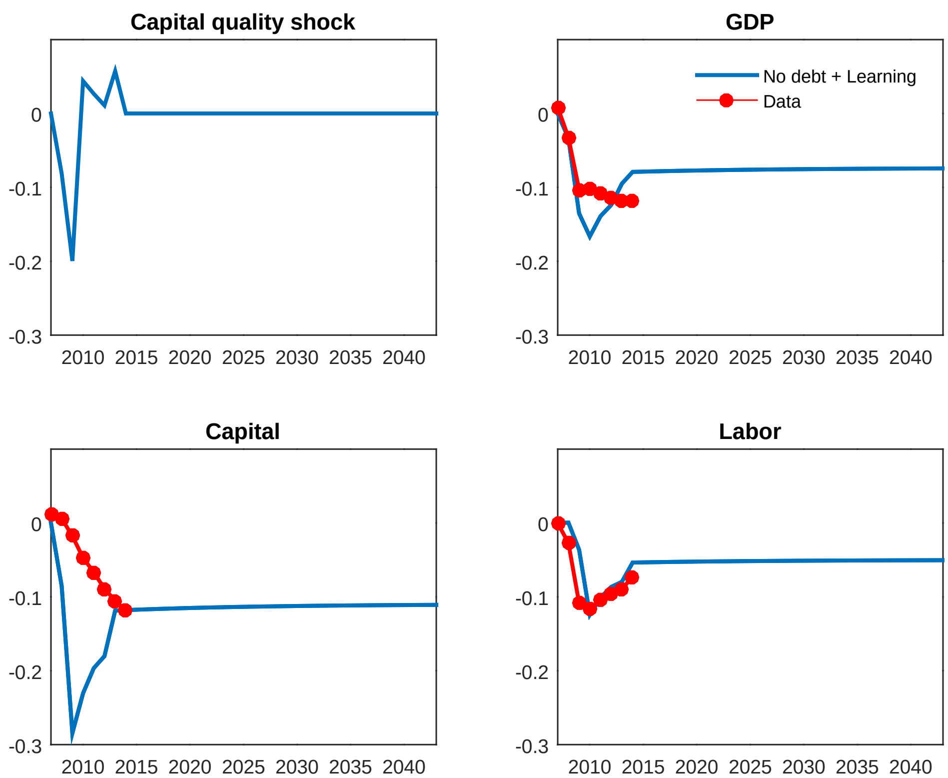

Figure 12: **Model vs data from 2008-2014.** Solid line is the baseline model subjected to the observed sequence of shocks from 2008-2014. The red circles are US data, in deviations from their pre-crisis trends.