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

Julian Kozlowski          Laura Veldkamp          Venky Venkateswaran*

NYU          NYU Stern, CEPR, NBER          NYU Stern, NBER

November 5, 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. Since changes in beliefs endure 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 study a production economy with agents who use standard econometric tools to estimate the distribution of aggregate shocks. When they observe a new shock, they re-estimate the distribution from which it was drawn. Even transitory shocks have persistent effects because, once observed, they stay forever in the agents' data set. We feed a time-series of US macro data into our model and show that our belief revision mechanism can explain the 12% downward shift in US trend output.

JEL Classifications: D84, E32

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

# 1   Introduction

The "Great Recession" was a deep downturn with long-lasting effects on credit markets, employment and output. As Figure 1 shows, post-WWII recessions in the US have typically had a distinct trough, followed by a sharp rebound toward a very stable trend line. In contrast, the Great Recession looks like a permanent level shift. We explore a simple explanation: This recession has been more persistent than others because it was perceived as an extremely unlikely event. Observing the crisis in 2008-09 caused agents to re-estimate 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 (defined as the probability of large adverse shocks).
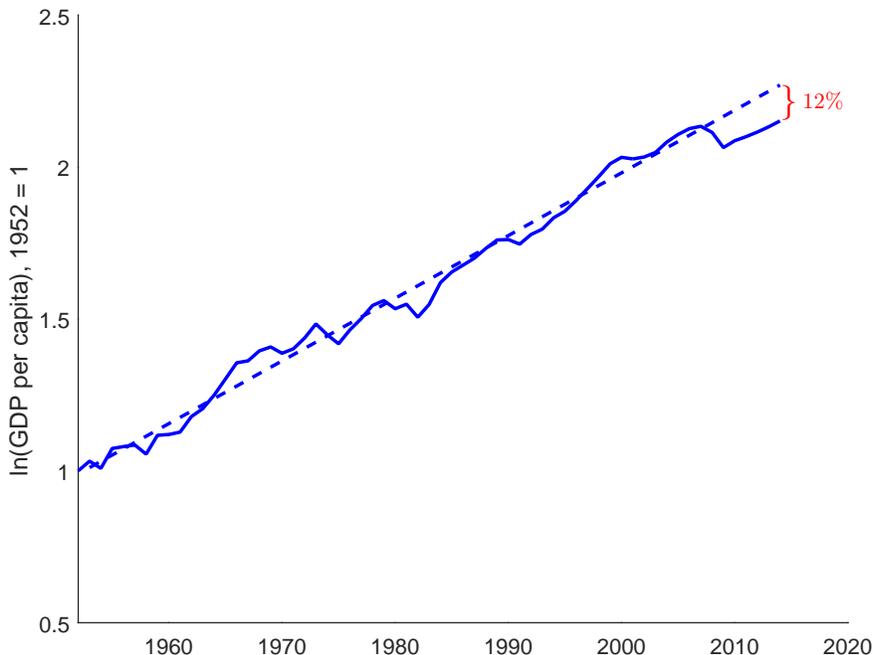


Figure 1: **Real GDP in the U.S. and its trend.**
*Dashed line is a linear trend that fits data from 1950-2007. In 2014, real GDP was 0.12 log points below trend.*

To measure how much tail risk rose, explain why it remained elevated, and explore its economic consequences, we analyze a production economy with agents who use standard econometric tools to estimate the distribution of aggregate shocks, in a flexible, 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 remain forever in the agents' data set. We feed a time-series of actual macro data for the post-war 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 12% downward shift in trend output in the aftermath of the Great Recession.

For most macro models, including belief-driven ones, producing persistent deviations from trend requires sufficiently persistent shocks.[1] With an exogenous change in persistence, such a theory cannot explain *why* this recession looked so different.

We propose a new type of belief-driven business cycles where persistence is endogenous and state-dependent. The key difference is that our agents are learning about a distribution, instead of about a hidden current or future state. In other words, fluctuations are persistent not because agents fear they are still in a "bad state," but rather, because the experience permanently changes their assessment of risk. So, why do some downturns appear to have more persistent effects? Extreme events, like the recent crisis, are rare and therefore, lead to significant changes in beliefs and long-run outcomes. Milder downturns, on the other hand, show up relatively more frequently in the agents' data set and therefore, have only small effects. In other words, while all changes to beliefs are permanent, deeper recessions induce larger, permanent belief changes.

This view is quantitatively successful in explaining the observed patterns in macro aggregates, in line with options data, consistent with popular narratives, and without a counterfactual sustained high credit spread. In Figure 2, the SKEW index, an option-implied measure of tail risk in equity markets, shows a clear rise since the financial crisis, with no subsequent decline.[2] Furthermore, popular narratives about the stagnation emphasize a change in "attitudes" or "confidence," that we capture with belief changes, and the reductions in debt financing that result:

> [Y]ears after U.S. investment bank Lehman Brothers collapsed, triggering a global financial crisis and shattering confidence worldwide, ... "The attitude toward risk is permanently reset." A flight to safety on such a global scale is unprecedented since the end of World War II. The implications are huge: Shunning debt... can starve the global economy. (Huffington Post Oct.6, 2013)

To make the case that realistic belief changes can explain stagnation, we augment an off-the-shelf economic model of debt-financed firms with our belief-formation mechanism. We then feed the calibrated model with post-war data through the recent financial crisis and compare its

---

[1]See, for example, Angeletos and La'O (2013) and Maćkowiak and Wiederholt (2010). Backus et al. (2015) analyze propagation in business cycle models.

[2]This index is constructed from market prices of out-of-the-money put options on the S&P 500. See http://www.cboe.com/micro/skew/introduction.aspx. Note that this is different from the VIX, which measures implied volatility of options struck at-the-money. The VIX rose dramatically at the onset of the crisis, but came down quite sharply afterwards.
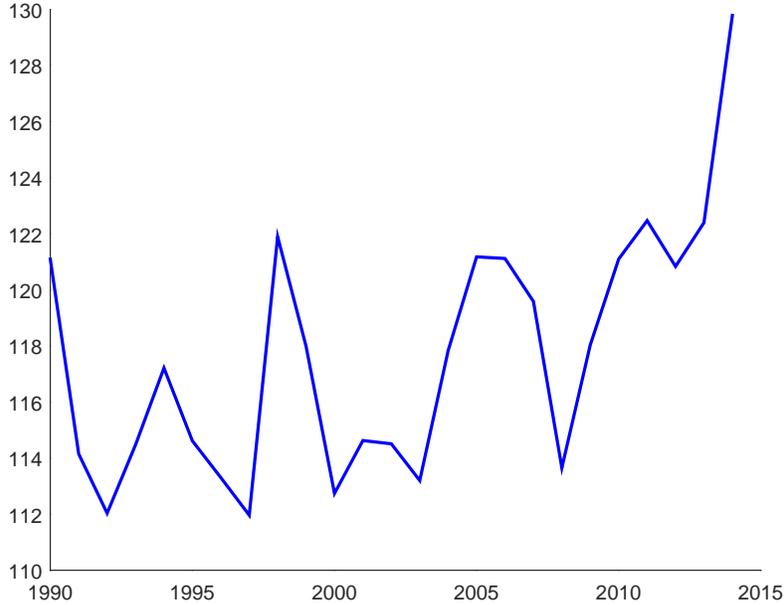
Figure 2: **The SKEW Index.**

*A measure of the market price of tail risk on the S&P 500, constructed using option prices. Source: Chicago Board Options Exchange (CBOE). 1990:2014.*

predictions to observed path of aggregate output. The novel feature of our analysis is the belief-formation mechanism, which is the source of endogenous persistence and which ties beliefs to observable data, avoiding free parameters. Intentionally, we do not add other innovative features to the economic environment and focus on developing a simple-to-execute and flexible toolkit to embed belief formation into standard quantitative macro models.

We begin in Section 2 by examining the belief-formation mechanism in isolation. We construct a time series of returns to US business capital and use it to show how our non-parametric estimation works. Agents estimate the distribution of aggregate capital returns by fitting a Gaussian kernel density function to real-time data. The extreme negative realizations during the financial crisis have a large and long-lived effect on beliefs. The theoretical underpinning of the persistence is the martingale property of beliefs: Future estimates of probabilities are, on average, equal to current ones. The intuition is simple: if an agent expects probabilities to predictably rise or fall in the future, then she should revise her current estimate. Of course, beliefs do change. If, for example, agents observe a large negative shock and then a sequence of moderate or positive shocks, the effect of the negative shock will diminish over time. But, averaging across all such possible paths for future beliefs yields the current estimated distribution.

To gauge the macroeconomic effects of this belief mechanism, Section 3 embeds it in a quantitative business cycle model. One approach would be to use a variant of a King et al. (1988) economy and allow agents to learn about aggregate technology (TFP) shocks. In such

4

an environment, belief revisions would be a source of endogenous persistence. However, this approach has two problems. First, productivity did not fluctuate much during the crisis. So belief revisions would be tiny. Second, technology shocks have little impact on investment. Even a surge in the risk of adverse TFP shocks barely changes the conditional distribution of capital returns. Therefore, it doesn't change the incentives to invest. Since most of secular stagnation shows up in depressed investment (Hall, 2014), a productivity-based model is not a promising explanation. If we want to have any shot of explaining the Great Recession, we need a model with an aggregate shock capable of generating large fluctuations in investment, in response to changes in tail risk. The state of the art in such models is Gourio (2012, 2013). Following other papers in the financial crisis literature, he uses aggregate shocks to capital quality (or equivalently, shocks to the effective depreciation rate of capital) as the driving force of aggregate fluctuations. Our contribution to this model is to assume that agents know that the shocks are i.i.d but do not know the distribution they are drawn from. This allows us to sharply highlight the ability of learning to generate endogenous persistence.

The other key ingredient is debt financing, which yields a tax advantage to the firm, but also leads to bankruptcy costs in the event of default. 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, financing investment with debt becomes less attractive. As a result, an economy with debt experiences an even larger drop in long-run investment and output than one without debt. We show that this amplification mechanism interacts with the size of the new shocks observed by the agent. The amplification is very powerful for large shocks (which, in turn, induce larger changes to beliefs), but almost negligible if the new data is closer to the mean. This accentuates the difference between extreme recessions and their milder counterparts.

Section 5 presents our quantitative results. First, 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 reports. We then use this to construct a time series of beliefs. Our calibrated model predicts that the increase in tail risk following the financial crisis triggers a cumulative drop in capital, employment and output of 17%, 8% and 12% respectively, with almost no rebound to trend. These predicted effects line up very well with the path of deviations from trend of the corresponding aggregate variables in the US data, despite the fact that (i) the shock itself is entirely transitory and (ii) the calibration and measurement of the shock does not make use of these aggregate data series directly.

Hall (2015) argues that credit market data is inconsistent with a beliefs-based explanation for stagnation. He notes that credit spreads rose only temporarily during the financial crisis and argues that the subsequent recovery to historically normal levels by 2010 implies a restoration

of confidence. This argument does not account for the decline in investment and borrowing as beliefs become more pessimistic, which offsets the increase in spreads. Our model teaches us that a surge in tail risk is consistent with the observed modest increase in spreads (about 4 bps relative to pre-crisis in the calibrated model). In other words, our belief mechanism can have significant and long-lived effects on economic activity, even if the effect on equilibrium credit spreads almost disappears.

Belief revisions are key to the model's ability to match the data - when we turn it off, the initial impact of the shock is similar, but all macro aggregates immediately start to rebound and ultimately return to the pre-crisis levels. In other words, without changes to beliefs, the effect of the financial crisis and the Great Recession are transitory. Similarly, a smaller shock, like the 1 standard deviation event which triggered the 2001 recession, generates fluctuations that are both smaller and less persistent. This demonstrates that the persistence of a shock's effect on economic activity depends on how rare it is. Events that are unlike previously-observed episodes produce larger belief revisions and therefore more persistent output fluctuations.

To better understand these results, we decompose the effects with a combination of numerical and analytical results. We analyze separately the role of changes in the mean (or the average shock), risk aversion, and debt. We find that each of these components contributes a significant proportion of the long-run decline in output.

Of course, the recent financial crisis was not the first time in history that the U.S. economy experienced large adverse economic shocks. The absences of the Great Depression in our agents' information set raises the question: How would access to more data, with large adverse shocks, affects the response of beliefs to the recent financial crisis? Section 5.2 addresses this issue using a simple counter-factual experiment - double the sample size and include another large financial crisis in that sample. We find that this attenuates the long-run effect of the recent crisis, by less than 30%. If agents do not place data from the distant past on an equal footing with more recent observation and instead, weight recent data more (1% annually), the predictions with the longer data sample are indistinguishable from the original model.

Not only does econometric belief formation add persistence to business cycle models, it also provides a new perspective on what it means for a price or quantity to be information insensitive. Debt is typically thought of as information-insensitive because its payoffs are flat throughout most of the state space. The only relevant region of the state space for debt payoffs is the left tail, where default occurs. Our learning mechanism reveals that the estimated probability of such tail events are particularly sensitive to new data because data on tail events is scarce. Thus, debt is very sensitive to beliefs exactly in a region where beliefs themselves are very sensitive to new information. This combination makes the value of debt and thus real economic activity quite sensitive to any new events that trigger revisions in tail probabilities.

6

Finally, our paper offers a broader methodological contribution. Obviously, no one truly knows the distribution of economic shocks. That is a useful fiction economists use to discipline beliefs. Our approach, on the other hand, assumes that agents learn from observable macro data using standard econometric tools, which imposes just as much discipline as more standard approaches. In addition, it is simple to execute and can be easily combined with a variety of sophisticated, quantitative macro models.

**Comparison to the literature**  Our framework, especially on the production side, builds on existing work on the macro consequences of exogenous shocks to disaster risks, e.g. Gourio (2012, 2013) and shocks to beliefs.[3] Our contribution, allowing agents to continually re-estimate the distribution of shocks, offers two key advantages. First, our belief revisions are not exogenous, but are firmly tied to observable data. Without discipline on the possible time-series of beliefs, many macroeconomic outcomes are rationalizable by the right signal and state process. 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 underlying states are persistent. Our mechanism is thus able to deliver persistent effects from transitory shocks and help explain 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. These uncertainty-based explanations leave two questions unanswered. First, why did the depressed level of economic activity continue long after the VIX and other measures of uncertainty had recovered? Our theory emphasizes tail risk. The SKEW index data in Figure 2 reveal that tail risk has lingered, making it a better candidate for explaining continued depressed output. Second, why were credit markets hardest hit and credit volume most persistently impaired after the crisis? Rises in tail risk hit the volume of credit because default is particularly sensitive to tail events.

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, these papers focus on endowment economies and do not analyze the potential for persistent effects in a setting with production. Pintus and Suda (2015) embed parameter learning in a production economy, but feed in persistent leverage shocks and explore the poten-

---

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

tial for amplification when agents hold erroneous initial beliefs about persistence. In Moriera and Savov (2015), learning changes demand for shadow banking (debt) assets. But, again, agents are learning about a hidden two-state Markov process, which has persistence built in. We, on the other hand, have transitory 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, 2013). 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 relies on hard-wired persistence in order to explain the prolonged stagnation. Our model aims to complement these theories by describing a simple mechanism that delivers endogenous persistence, explains why extreme events, like the recent crisis, lead to more persistent responses than milder downturns, and is easy to implement and combine with many existing economic frameworks.

Finally, our paper contributes to the recent literature on secular stagnation. Eggertsson and Mehrotra (2014) argue that a combination of low effective demand and the zero lower bound on nominal rates can generate a long-lived slump. In contrast, Gordon (2014), Anzoategui et al. (2015) and others attribute stagnation to a decline in productivity. Hall (2015) surveys other theories as well. While these alternatives may well be part of the explanation, we offer a simple mechanism that can reconcile the recent stagnation with the propagation of business cycle shocks in normal times as well.

The rest of the paper is organized as follows. Section 2 describes the belief-formation mechanism. Section 3 presents the economic model. Section 4 shows the measurement of shocks and calibration of the model. Section 5 analyzes the main results of the paper while Section 6 decomposes the principal economic forces driving the results. Finally, Section 7 concludes. Details are gathered in Appendix A and B.

# 2  Belief Formation

The main contribution of this paper is to explain why tail risk fluctuates and why such fluctuations can create persistent responses to transitory shocks. Before laying out the whole model, we begin by explaining the novel part of the paper – how agents form beliefs and why they are persistent. These results are based on statistical properties of beliefs and therefore, are more general than the results of the specific economic model in the following section, which is used to quantify the effect of belief changes in the context of the US economy.

No one knows the true distribution of shocks to the economy. We estimate such distributions, updating our beliefs as new data arrives. The first step is to choose a particular estimation

procedure. A common approach is to assume a normal distribution and estimate its parameters (namely, mean and variance). While tractable, this has the disadvantage that the normal distribution, with its thin tails, is unsuited to thinking about changes in tail risk. We could choose a distribution with more flexibility in higher moments. However, this will raise obvious concerns about the sensitivity of results to the specific distributional assumption used. To minimize such concerns, we take a non-parametric approach and let the data inform the shape of the distribution.

Specifically, we employ a kernel density estimation procedure, one of most common approaches in non-parametric estimation. Essentially, it approximates the true distribution function with a smoothed version of the empirical distribution. By using the widely-used normal kernel, we impose a lot of discipline on our learning problem but also allow for considerable flexibility. We also studied a handful of other kernel specifications, which yielded similar results.[4]

Consider an aggregate shock $\phi_t$ whose true density $g$ is unknown to agents in the economy. The agents do know that the shock $\phi_t$ is i.i.d.. Assuming independence serves a useful purpose: It ensures that the persistence in long run outcomes comes only from our mechanism, changes in beliefs. While persistent shocks are also possible and tractable, a theory that assumes persistence cannot answer our main question: why some aggregate fluctuations are more persistent than others.

The information set of our agents at time $t$, denoted $\mathcal{I}_t$, includes the history of all shocks $\phi_t$ observed up to and including $t$. They use this available data to construct an estimate $\hat{g}_t$ of the true density $g$. Formally, at every date, agents construct the following normal kernel density estimator of the pdf $g$

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

where $\Omega(\cdot)$ is the standard normal density function, $\kappa_t$ is the smoothing or bandwidth parameter and $n_t$ is the number of available observations of at date $t$. As new data arrives, agents add the new observation to their data set and update their estimates, generating a sequence of beliefs $\{\hat{g}_t\}$.

The key mechanism in the paper is the persistence of belief changes induced by purely transitory $\phi_t$ shocks. This stems from the so-called martingale property of beliefs - i.e. conditional on time-$t$ information $(\mathcal{I}_t)$, the estimated distribution is a martingale. Thus, on average, the

---

[4]Kernels we explored included other non-parametric kernels like Epinechnikov, kernels designed to better capture tail risk like Champernowne, as well as semi-parametric kernels with Pareto tails and the g-and-h family which covers several transformations of the normal distribution. Each alternative yielded similar economic predictions because new data increased the tail probabilities of each distribution in a similar way. For a detailed discussion of nonparametric estimation, see Hansen (2015).

agent expects her future belief to be the same as her current beliefs. This property holds exactly if the bandwidth parameter $\kappa_t$ is set to zero[5]. In our empirical implementation, in line with the literature on non-parametric assumption, we use the optimal bandwidth[6]. This leads to a smoother density but also means that the martingale property does not hold exactly. We will however show, numerically, that the deviation are minuscule, both for the illustrative example in this section and in our full model. In other words, the kernel density estimator with the optimal bandwidth is, for all practical purposes, a martingale $\mathbb{E}_t\left[\hat{g}_{t+j}\left(\phi\right)|\mathcal{I}_t\right] \approx \hat{g}_t\left(\phi\right)$. As a result, any changes in beliefs induced by new information are, in expectation, permanent. This property plays a central role in generating long-lived effects from transitory shocks.

**Example: Capital returns**    We now illustrate how this belief formation mechanism works by applying the estimation procedure described above to the time series of returns on business (i.e. non-residential) capital for the US economy in the post-WWII era. We picked this time series in part to the striking effects of the financial crisis in 2008-'09 on this variable. For the purposes of this section, we treat returns as exogenous. In our economic model in the next section, they will emerge endogenously from the interaction of more primitive shocks and the economic environment. Our analysis here previews the dynamics of beliefs in that more complete setup.

**Measuring returns**    We begin with the standard definition of returns

$$r_t = \frac{d_t}{p_{t-1}} + \frac{p_t}{p_{t-1}} \tag{1}$$

where the two components denote the dividend yield and capital gains in $t$. For the former, we draw on the work of Gomme et al. (2011), who use data from the National Income and Product Accounts (NIPA) with careful adjustments for proprietors' income, rents and taxes to construct a measure of payments to capital. We use the same data as in that paper and follow their methodology to generate an annual time series of the dividend yield component (defined

---

[5]As $\kappa_t \to 0$, the CDF of the kernel converges to $\hat{G}_t^0\left(\phi\right) = \frac{1}{n_t}\sum_{s=0}^{n_t-1}\mathbf{1}\left\{\phi_{t-s} \leq \phi\right\}$. Then, for any $\phi$, $j \geq 1$

$$\mathbb{E}_t\left[\hat{G}_{t+j}^0\left(\phi\right)\Big|\mathcal{I}_t\right] = \mathbb{E}_t\left[\frac{1}{n_t+j}\sum_{s=0}^{n_t+j-1}\mathbf{1}\left\{\phi_{t+j-s} \leq \phi\right\}\Bigg|\mathcal{I}_t\right]$$

$$\mathbb{E}_t\left[\hat{G}_{t+j}^0\left(\phi\right)\Big|\mathcal{I}_t\right] = \frac{n_t}{n_t+j}\hat{G}_t^0\left(\phi\right) + \frac{j}{n_t+j}\mathbb{E}_t\left[\mathbf{1}\left\{\phi_{t+1} \leq \phi\right\}|\mathcal{I}_t\right]$$

Thus, future beliefs are, in expectation, a weighted average of two terms - the current belief and the distribution from which the new draws are made. Since our best estimate for the latter is the current belief, the two terms are exactly equal, implying $\mathbb{E}_t\left[\hat{G}_{t+j}^0\left(\phi\right)\Big|\mathcal{I}_t\right] = \hat{G}_t^0\left(\phi\right)$.

[6]See Hansen (2015).

as pre-tax capital income divided by the value of capital stock) for the period from 1948-2009.[7] For the capital gain component, we use annual data on non-financial assets for US corporate entities. 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 combine the two series, denoted $NFA_t^{HC}$ and $NFA_t^{RC}$ respectively and construct the following measure of capital gains in period $t$

$$\left(\frac{p_t}{p_{t-1}}\right)^{Nominal} = \frac{NFA_t^{RC}}{(1-\delta)NFA_{t-1}^{RC} + \text{Invt}_t} \tag{2}$$

$$= \frac{NFA_t^{RC}}{(1-\delta)NFA_{t-1}^{RC} + NFA_t^{HC} - (1-\delta)NFA_{t-1}^{HC}} \tag{3}$$

The formula in (3) uses changes in the market/replacement value series, adjusted for new investments to isolate the part arising from revaluation effects. We then deflate this series using an aggregate price index to remove purely nominal changes and add it to the dividend yield described earlier. The combined return series is plotted in the first panel of Figure 3. It shows that realized returns during the financial crisis were significantly lower than any that were observed throughout the entire sample.
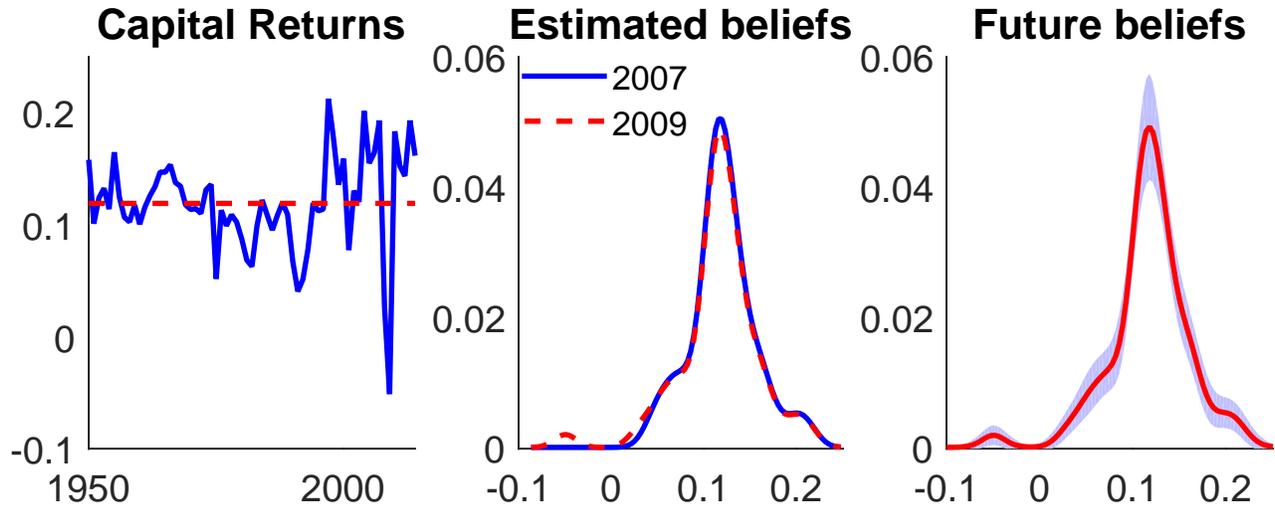


Figure 3: **Estimated Beliefs about Capital Returns.**
*The first panel shows the realized capital returns. The second panel shows the estimated kernel density for 2007 and 2009. The third panel shows the mean belief (along with a 2 standard deviation band) in 2039 (computed by simulating data for the period 2010-2039 using the estimated distribution in 2009).*

---

[7]We thank the authors for sharing their data and code with us.

**Estimated belief changes** The estimated distributions using this data for two dates - 2007 (pre-crisis) and 2009 (post-crisis) - are shown in the second panel of Figure 3. We note that these adverse realizations lead to an increase in tail risk. The 2009 distribution, $\hat{g}_{2009}$ shows a pronounced hump in the density around the 2008 and 2009 realizations, relative to the pre-crisis one. Crucially, even though these negative realizations were short-lived, this increase in left tail risk persists. To see how persistent they are, we ask the following question: What would be the mean belief in 2039 if the economy continued to be hit by shocks drawn from the estimated distribution in 2009 and beliefs were re-estimated using the combined actual and simulated data set? The third panel of Figure 3 shows the results from Monte Carlo simulations. Obviously, each simulated path gives rise to a different estimate, but averaging across all those paths yields the 2009 distribution. This simulation illustrates how tail risk induced by financial crisis may never go away. This pattern is reminiscent of Figure 2, which showed that price of tail risk in equity options markets continues to remain high.

Note that the left tail "hump" never disappears because we are drawing from the $\hat{g}_{2009}$ distribution. If we drew future data from a distribution without tail risk (e.g. $\hat{g}_{2007}$), the hump would shrink over time. This raises the question of why we simulate drawing from $\hat{g}_{2009}$. Given that $\hat{g}_{2009}$ is the agents' best estimate of the true distribution, drawing from some other distribution would amount to assuming that we, as modelers, have more information than the agents in our model, which is contrary to the spirit of our exercise. Therefore, throughout this paper, we endow agents with the same information that we ourselves have.

Thus, every new shock to capital returns ($\phi_t$), even ones that are transitory, has a permanent expected effect on beliefs. To assess the implications of these belief changes for macroeconomic outcomes, we need a model that maps shocks and beliefs into investment, hiring and production decisions. The model laid out in the following section does just that - it is a set of specific assumptions that highlight the effect of tail risk on aggregate variables. However, we wish to re-iterate that this flexible, non-parametric approach to belief formation is a simple tool that could be embedded in many other macroeconomic models and holds the potential for persistent responses to transitory shocks.

# 3   Economic Model

This section describes the setup and equilibrium of a production economy. Since our innovation - and primary contribution - is to introduce real-time estimation of beliefs as a source of endogenous persistence, we do not want to introduce additional new features of the economic environment. Therefore, we use an existing model of the economy, Gourio (2013), which in turn builds on Gertler and Karadi (2011) and Gourio (2012). In the following sections, we use

this existing framework to quantify the size of the responses in capital, output and labor to our estimated changes in beliefs.

## 3.1 Setup

This section simply describes a version of Gourio (2013) model, where firms finance investment and payroll expenses using a combination of debt and equity financing and are subject to aggregate and idiosyncratic shocks.

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

$$U_t = \left[ (1 - \beta) \left( C_t - \zeta \frac{L_t^{1+\gamma}}{1 + \gamma} \right)^{1-\psi} + \beta E_t \left( U_{t+1}^{1-\eta} \right)^{\frac{1-\psi}{1-\eta}} \right]^{\frac{1}{1-\psi}} \tag{4}$$

where $\psi$ is the inverse of the intertemporal elasticity of substitution, $\eta$ indexes risk-aversion and $\gamma$ 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 $A k_{it}^{\alpha} l_{it}^{1-\alpha}$, 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 $\phi_t$. A firm that enters the period with capital $\hat{k}_{it}$ and is hit by a shock $\phi_t$ has effective capital $k_{it} = \phi_t \hat{k}_{it}$. These capital quality shocks $\phi_t \sim g(\cdot)$ are the only aggregate disturbances in our economy.

These shocks scale up or down the effective capital stock. Of course, they are not to be interpreted literally - it is hard to visualize shocks that regularly wipe out fractions of the capital or create it out of thin air. Instead, these shocks are a simple, yet imperfect way to model permanent changes in the economic value of a piece of capital. It allows us to capture the idea that a hotel built in 2007 in Las Vegas may still be standing, but may deliver much less economic value after the financial crisis. Capital quality shocks have been employed for a similar purpose in Gourio (2012), as well as in a number of recent papers on financial frictions, crises and the Great Recession (e.g., Gertler et al. (2010), Gertler and Karadi (2011), Brunnermeier and Sannikov (2014)). Their use in macroeconomics and finance, however, goes back at least to Merton (1973), who uses them to reconcile the high volatility in asset returns with less-volatile macro series. Until we have a firmer understanding of why returns are so volatile, we need to use stand-in shocks like these if we want this kind of model to have any shot of speaking to both financial and macro data.

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)

$$\Pi_{it} = v_{it} \left[ A k_{it}^{\alpha} l_{it}^{1-\alpha} + (1 - \delta) k_{it} \right] \tag{5}$$

where $\delta$ 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, $F$.[8] The mean of the idiosyncratic shock is normalized to be one: $\int v_{it} \; di = 1$. The primary role of these shocks is to induce an interior default rate in equilibrium, allowing a more realistic calibration.

**Labor, credit markets and default:** We make two additional assumptions about labor markets. First, firms hire labor in advance, i.e. before observing the realizations of aggregate and idiosyncratic shocks. Second, wages are non-contingent - in other words, workers are promised a non-contingent payment and face default risk. These features serve to create an additional source of leverage, but are not central to our results, whether qualitatively and quantitatively.

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 period $t + 1$ 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, and only if,

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

where $\Gamma_{t+1}$ 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 ($\Pi_{it+1}$) and the *total* obligations of the firm to both bondholders 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

---

[8]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.

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.[9]

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 $q\left(\cdot\right)$ 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,

$$q\left(\hat{k}_{it+1}, l_{it+1}, B_{it+1}, S_t\right) = \mathbb{E}_t M_{t+1} \left[ r_{it+1} + (1 - r_{it+1}) \frac{\theta \tilde{V}\left(\Pi_{it+1}, S_{t+1}\right)}{B_{it+1}} \right] \tag{6}$$

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

$$M_{t+1} = \left(\frac{dU_t}{dC_t}\right)^{-1} \frac{dU_t}{dC_{t+1}} = \beta \left[E_t\left(U_{t+1}^{1-\eta}\right)\right]^{\frac{\eta-\psi}{1-\eta}} U_{t+1}^{\psi-\eta} \left(\frac{u\left(C_{t+1}, L_{t+1}\right)}{u\left(C_t, L_t\right)}\right)^{-\psi}. \tag{7}$$

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 $\Pi_t$, 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[10] $q_{it}$ 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.[11]

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

---

[9]Note also that this means that default does not destroy resources - the penalty is purely private. This is not crucial - it is straightforward to relax this assumption by assuming that all or part of the cost of the default represents physical destruction of resources.

[10]In a slight abuse of notation, we denote the equilibrium bond price by $q_{it}$.

[11]The subsidy is assumed to be paid by a government that finances it through a lump-sum tax on the representative household.

capital stock (the undepreciated current capital stock is included in $\Pi_{it}$), plus the revenue earned from issuing new debt, including its tax subsidy

$$d_{it} = \Pi_{it} - B_{it} - \hat{k}_{it+1} + \chi q_{it} b_{it+1}. \tag{8}$$

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

As mentioned earlier, workers are subject to the risk of default. Since this risk is the same as that faced by bondholders, the expected discounted value of a wage promise $w_{it+1}$ is obtained by simply multiplying it by the equilibrium bond price

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

where the expectation is taken over aggregate and idiosyncratic shocks. In other words, the workers are essentially paid through bonds. Note that workers are also members of the representative family and therefore, evaluate their wage claims using the stochastic discount factor, $M_{t+1}$.

From the household's problem, we can derive the following optimality condition for labor supply

$$
\begin{aligned}
w_{it+1}q_{it}\frac{dU_t}{dC_t} &= \frac{dU_t}{dL_{t+1}} \\
w_{it+1}q_{it} &= \left(\frac{dU_t}{dC_t}\right)^{-1}\frac{dU_t}{dL_{t+1}} \equiv \mathcal{W}_t
\end{aligned}
\tag{9}
$$

The expected wage rate, weighted by the economy-wide stochastic discount factor $M_{t+1}$ for every firm in the economy must equal the marginal rate of substitution of the representative household (denoted $\mathcal{W}_t$) and is obtained by multiplying the promised wage, $w_{it+1}$, by the bond price $q_{it}$.

**Timing and value functions:** 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 $\phi_t$ 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.

In recursive form, the problem of the firm is

$$V\left(\Pi_{it}, B_{it}, S_t\right) = \max\left[0, \max_{d_{it}, \hat{k}_{it+1}, b_{it+1}, w_{it+1}, l_{it+1}} d_{it} + \mathbb{E}_t M_{t+1} V\left(\Pi_{it+1}, B_{it+1}, S_{t+1}\right)\right] \qquad (10)$$

subject to

$$\text{Dividends:} \quad d_{it} \leq \Pi_{it} - B_{it} - \hat{k}_{it+1} + \chi q_{it} b_{it+1}$$

$$\text{Discounted wages:} \quad \mathcal{W}_t \leq w_{it+1} q\left(\hat{k}_{it+1}, l_{it+1}, B_{it+1}, S_t\right)$$

$$\text{Future obligations:} \quad B_{it+1} = b_{it+1} + w_{it+1} l_{it+1}$$

$$\text{Resources:} \quad \Pi_{it+1} = v_{it+1}\left[A(\phi_{t+1}\hat{k}_{it+1})^\alpha l_{it+1}^{1-\alpha} + (1-\delta)\phi_{t+1}\hat{k}_{it+1}\right]$$

$$\text{Bond price:} \quad q\left(\hat{k}_{it+1}, l_{it+1}, B_{it+1}, S_t\right) = \mathbb{E}_t M_{t+1}\left[r_{it+1} + (1 - r_{it+1})\frac{\theta \tilde{V}_{it+1}}{B_{it+1}}\right]$$

The first max operator in (10) 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 aggregate shock distribution in $S_t$ as given. 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 three aggregate objects that are relevant for the firm's problem are the wage rate $\mathcal{W}_t$, the stochastic discount factor, $M_{t+1}$ and beliefs. These depend on aggregate consumption and labor. Therefore, the aggregate state $S_t$ consists of $(\Pi_t, L_t, \mathcal{I}_{it})$ where $\Pi_t \equiv AK_t^\alpha L_t^{1-\alpha} + (1-\delta)K_t$ is the aggregate resources available.

**Information, beliefs and equilibrium** The economy-wide information set, denoted $\mathcal{I}_t$, includes the history of all shocks $\phi_t$ observed up to and including time-$t$. For now, we specify a general function, denoted $\Psi$, which maps $\mathcal{I}_t$ into an appropriate probability space. The expectation operator $\mathbb{E}_t$ is defined with respect to this space. In the following section, we make this more concrete using the kernel density estimation procedure outlined in section 2 to map the information set into beliefs.

For a given belief function $\Psi$, a recursive equilibrium is a set of (i) functions for aggregate

consumption and labor that maximize (4) subject to a budget constraint, (ii) firm value functions and associated policy functions that solve (10) , taking the bond price, stochastic discount factor and wage functions - (6), (7) and (9) respectively - as given. (iii) aggregate consumption and labor are consistent with individual choices.

## 3.2  Characterization and solution

The equilibrium of the economic model is a solution to a set of non-linear equations. In this section, we derive those equations and explore how tail events and the subsequent changes in beliefs affect them. In particular, how tail risk changes the incentives to hire and invest. We present only the key equations here and relegate the detailed derivations to Appendix A.

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

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

In Appendix A, we show that the optimality condition with respect to capital can be expressed as follows[12]

$$
1 + \chi \mathcal{W}_t \frac{l_{t+1}}{\hat{k}_{t+1}} = \mathbb{E}[M_{t+1} R_{t+1}^k] + (\chi - 1)\frac{B_{it+1}}{\hat{k}_{it+1}}q_t - (1 - \theta)\mathbb{E}[M_{t+1}R_{t+1}^k h(\underline{v})] \tag{11}
$$

where

$$
R_{t+1}^k = \frac{A\phi_{t+1}^\alpha \hat{k}_{t+1}^\alpha l_{t+1}^{1-\alpha} + (1 - \delta)\phi_{t+1}\hat{k}_{t+1}}{\hat{k}_{t+1}}
$$

$$
h(\underline{v}) \equiv \int_{-\infty}^{\underline{v}} v f(v)dv.
$$

The term $R_{t+1}^k$ is the *ex-post* per-unit, pre-wage return on capital, while $h(\underline{v})$ is the default-weighted expected value of the idiosyncratic shock.[13]

The first term on the right hand side of (11) is the usual expected direct return from investing, weighted by the stochastic discount factor. This would be the only term in a world without debt (or in the problem of a planner maximizing the representative household's utility). The other two terms are related to debt. The second term reflects the indirect benefit to investing arising from the tax advantage of debt - for each unit of capital, the firm raises $\frac{B_{it+1}}{\hat{k}_{it+1}}q_t$

---

[12]Since all firms are identical, they make symmetric choices and accordingly, we suppress the $i$ subscript.
[13]$h(\underline{v})$ can also be expressed as $\int_{-\infty}^{\infty} v(1 - r)f(v)dv$.

from the bond market and earns a subsidy of $\chi - 1$ on the proceeds. The last term is the cost of this strategy - default-related losses, equal to a fraction $1 - \theta$ of available resources.

All three components of the return to investing are influenced by - and therefore, change with - beliefs. When agents see large negative realizations, they update their beliefs about the likelihood of similar outcomes in the future. This increase in tail risk drives down the expected direct returns. It also raises risk premia (acting through the stochastic discount factor $M_{t+1}$). Both these effects serve to lower the first term, reducing incentives to invest. The effects on the debt-related terms in (11) is more subtle. For a given level of debt, more pessimistic beliefs increase the probability of default, reducing the bond price and the tax subsidy. However, in general, it could have ambiguous effects on the expected costs of default.[14] Numerically, we find that the change in beliefs post-2009 reduces the net returns to capital from the debt-related terms as well. Intuitively, tail risk lowers the effective tax subsidy (through lower bond prices) and raises default-related losses.

A few more lines of algebra yield the following representation of equation (11)

$$1 + \chi \mathcal{W}_t \frac{l_{t+1}}{\hat{k}_{t+1}} = \mathbb{E}_t \left[ M_{t+1} R_{t+1}^k J^k(\underline{v}) \right] \tag{12}$$

where

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

$J^k(\underline{v})$ reflects the net effect of distortions induced by debt and can be interpreted as a wedge, which distorts equilibrium capital choice away from the choices of a planner. In the absence of debt (e.g. if $\chi = 1$), $J^k(\underline{v}) = 1$, reducing (11) to a standard Euler equation.

The optimality condition for labor looks quite similar. As with capital, firms equate the expected marginal cost of an additional unit of labor, namely $\mathcal{W}_t$, with its expected marginal product, adjusted for the effect of additional promised wages on the cost of default

$$\chi \mathcal{W}_t = \mathbb{E}_t \left[ M_{t+1}(1 - \alpha) A\phi_{t+1}^\alpha \left( \frac{\hat{k}_{t+1}}{l_{t+1}} \right)^\alpha J^l(\underline{v}) \right] \tag{13}$$

where

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

Finally, rather than work with the choice of bonds, it turns out to be convenient to analyze the problem of choice of leverage (defined as the ratio of total obligations $B_{t+1}$ to capital $\hat{k}_{t+1}$).

---

[14]To see this, note that tail risk makes default more likely but also reduces the available resources and therefore, output losses, in the event of default.

The associated optimality condition is given by

$$(1 - \theta)\, \mathbb{E}_t\left[M_{t+1}\underline{v}f\left(\underline{v}\right)\right] = \left(\frac{\chi - 1}{\chi}\right)\mathbb{E}_t\left[M_{t+1}\left(1 - F\left(\underline{v}\right)\right)\right] \tag{14}$$

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, $(12)$ -$(14)$, along with those from the household side - in particular, the labor supply condition $(9)$ - characterize the equilibrium of this economy, for any belief formation process. In the following sections, we solve these equations numerically.

# 4 Measurement and Calibration

In this section, we describe how we use macro data to pin down beliefs. One of the key strengths of our belief-driven theory is that, by assuming that agents form beliefs as an econometrician would, we can use observable data to discipline beliefs. We also parameterize the model to match key features of the US economy and describe key aspects of our computational approach.

## 4.1 Measuring capital quality shocks

To construct a time series of $\phi_t$, we use the two series on non-financial assets in the US economy from the Flow of Funds reports that we used in Section 2. Recall that one series was evaluated at historical cost, $NFA_t^{HC}$, while the other at replacement cost or market value, $NFA_t^{RC}$. In the language of our model, the latter series can be interpreted as effective capital, or $K_t$. Letting $X_{t-1}$ denote investment in period $t - 1$ and $P_t$ the nominal price of capital goods in $t$, we can formally map these objects into their model counterparts as follows

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

To adjust for changes in the nominal price $P_t^k$, we use the price index for non-residential investment from the National Income and Product Accounts (denoted $PINDX_t$).[15] This allows

---

[15]Our results are robust to alternative measures of nominal price changes, e.g. computed from the price index for GDP or Personal Consumption Expenditure, see Appendix B.1.

us to recover the quality shock $\phi_t$

$$\phi_t = \frac{K_t}{\hat{K}_t} = \left(\frac{P_t^k K_t}{P_{t-1}^k \hat{K}_t}\right)\left(\frac{P_{t-1}^k}{P_t^k}\right)$$
$$= \left(\frac{NFA_t^{RC}}{(1-\delta)NFA_{t-1}^{RC} + NFA_t^{HC} - (1-\delta)NFA_{t-1}^{HC}}\right)\left(\frac{PINDX_{t-1}^k}{PINDX_t^k}\right) \qquad (15)$$

where the second line replaces $\frac{P_{t-1}^k}{P_t^k}$ with $\frac{PINDX_{t-1}^k}{PINDX_t^k}$.

Using the measurement equation (15), we construct an annual time series for capital quality shocks for the US economy since 1950. The left panel of Figure 4 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 these numbers in context, the mean and standard deviation of the series from 1950-2007 were 1 and 0.03 respectively.

We then apply our 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 panel of Figure 4. They show that the Great Recession induced a significant increase in the perceived likelihood of extreme negative shocks. The estimated density for 2007 implies almost zero mass below 0.90, while the one for 2009 attach a non-trivial (approximately 2.5%) probability to this region of the state space.
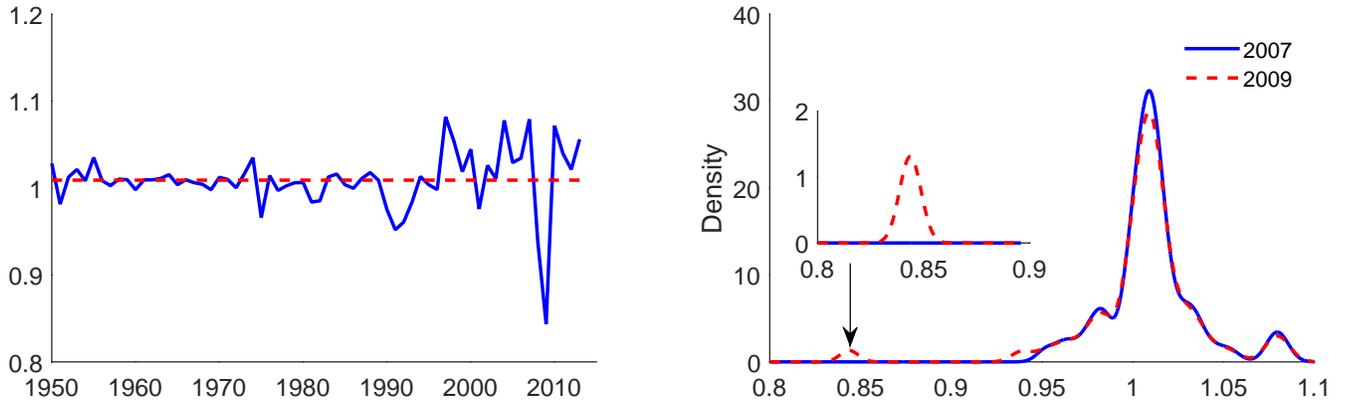


Figure 4: **Capital quality shocks.**
*The left panel shows the time series of $\phi_t$ measured from the US data using (15). The right panel shows the estimated kernel densities in 2007 (solid) and 2009 (dashed) respectively. The change in left tail shows the effect of the Great Recession.*

## 4.2  Calibration

A period is interpreted as a year. We choose the discount factor $\beta$ and depreciation $\delta$ to target a steady state capital-output ratio of 3.5 (this is taken from Cooley and Prescott (1995)) and an investment-output ratio of 0.12 (this is the average ratio of non-residential investment to output during 1950-2007 from NIPA accounts)[16]. The share of capital in the production, $\alpha$, is 0.40, which is also taken from Cooley and Prescott (1995). The recovery rate upon default, $\theta$, 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 of 0.02.[17] The labor supply parameter, $\gamma$, is set to 0.5, in line with Midrigan and Philippon (2011), corresponding to a Frisch elasticity of 2. The labor disutility parameter $\zeta$ 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 and set $\psi = 0.5$ (or equivalently, an IES of 2) and $\eta = 10$.[18] The tax advantage parameter $\chi$ is chosen to match a leverage target of 0.70, which is 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, about 0.5 in US data - from Gourio (2013)). Table 1 summarizes the resulting parameter choices.

## 4.3  Numerical solution method

To solve for the equilibrium in our baseline model, at each date, we need to compute current beliefs $\hat{g}_t$ as well as the probability distribution over future beliefs, namely over $\hat{g}_{t+1}$. Thus, one of the aggregate state variables is a distribution over a distribution. Carrying around this high-dimensional object makes computations almost unviable. Therefore, we approximate the solution by assuming that agents evaluate the future using the time-$t$ conditional mean of future beliefs, $\hat{g}_{t+s}$. By the martingale property of beliefs, we know that $\hat{g}_t$ is the mean of future belief distributions. In other words, for each date $t$, we replace the distribution of future probability distributions with $\hat{g}_t$. Thus, as in many other types of approximation routines, we are summarizing a distribution state variable with its mean. In Appendix B.2, we conduct some numerical experiments which show that our solution method delivers accurate results.

---

[16]This leads to values for $\beta$ and $\delta$ of 0.91 and 0.03 respectively. These are lower than other estimates in the literature. However, when we used an alternative calibration strategy with $\delta = 0.06$ (which is consistent with reported depreciation rates in the Flow of Funds data) and $\beta = 0.95$ (which leads to the same capital-output ratio), the resulting impulse responses were almost identical.

[17]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.

[18]See discussion in Gourio (2013).

| Parameter | Value | Description |
|---|---|---|
| **Preferences:** | | |
| $\beta$ | 0.91 | Discount factor |
| $\eta$ | 10 | Risk aversion |
| $\psi$ | 0.50 | 1/Intertemporal elasticity of substitution |
| $\gamma$ | 0.50 | 1/Frisch elasticity |
| $\zeta$ | 1 | Labor disutility |
| **Technology:** | | |
| $\alpha$ | 0.40 | Capital share |
| $\delta$ | 0.03 | Depreciation rate |
| $\hat{\sigma}$ | 0.25 | Idiosyncratic volatility |
| $A$ | 1 | TFP |
| **Debt:** | | |
| $\chi$ | 1.06 | Tax advantage of debt |
| $\theta$ | 0.70 | Recovery rate |

Table 1: **Parameters**

# 5   Main Results

Our main goal in this section is to quantify the size and persistence of the macroeconomic response to a large but transitory shock $\phi_t$ and explore how belief updating, shock size, and additional past data alter the economic responses in the long run.

To assess the economic effect of our revisions in beliefs, we begin by estimating $\hat{g}_{2007}$ using historical data on $\phi_t$ from 1950-2007, measured using the strategy outlined in Section 4 and computing the stochastic steady state associated with this distribution.[19] Then, starting from this steady state, 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, $\hat{g}_{2009}$. Then, we simulate future paths (drawing from this revised estimate $\hat{g}_{2009}$) and plot the mean future path of various aggregate variables in Figure 5. The top left panel shows the time path for $\phi_t$ (as deviations from its average value). The remaining panels show the behavior of output, investment and employment over time. They show a pattern of prolonged stagnation, where the economy (on average) never recovers from the negative shocks in 2008-'09 as all aggregate variables move towards the new, lower (stochastic) steady state.

It is important to note that the results in Figure 5 do not imply that stagnation will necessarily continue forever. It tells us that, from the perspective of an agent with the 2009 information set, recovery is not expected. If, for example, we simulated post-2009 paths drawing

---

[19]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.
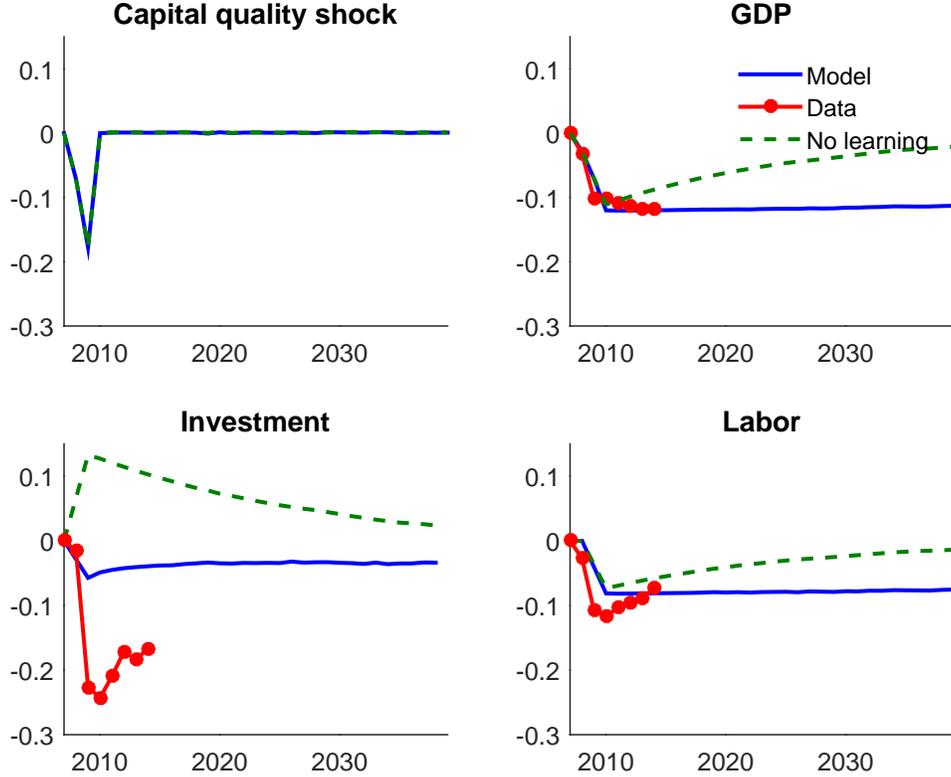
Figure 5: **Large negative shocks create extremely persistent responses in output, investment and labor.**

*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 2008-2014. Dashed line (no learning) is an identical model where agents believe that shocks are drawn from $\hat{g}_{2009}$ and never revise those beliefs.*

from a more optimistic distribution (i.e. with lower probability of crisis), the long-run outcomes may not be very different from the pre-crisis levels. In such a scenario, the effects of the Great Recession will still be persistent, because learning is relatively slow, but not permanent. However, such an exercise requires taking a stand on the objectively true distribution of shocks. The central premise of our analysis is that no one, neither agents nor model-builders, knows the true distribution. Therefore, we put ourselves on the same footing as the agents in our model and compute future expected outcomes from their perspective, using all information available. Since beliefs are martingales, changes in beliefs are expected to be permanent, and no recovery is anticipated.

The solid line with circles in Figure 5 plots the actual data (in deviations from their respective 1950-2007 trends) for the US economy.[20] As the graph shows, the model's predictions for

---

[20]Data on output and labor input are obtained from Fernald (2014). Data on investment is obtained from the series for non-residential investment reported by Bureau of Economic Analysis. Each series is detrended using a log-linear trend estimated using data from 1950-2007.

GDP and labor line up remarkably well with the recent data, though none of these series were used in the calibration or measurement of the aggregate shock $\phi_t$. The predictions for output, for example, are almost identical to the observed data, while the predicted path for employment seems to lag (and slightly underpredict) the actual changes.[21] In the data, employment dropped sharply in 2008-'09, almost contemporaneously with the negative shocks and then recovered slightly. In the model, however, the drop occurs later, largely due to the assumption that labor is chosen in advance.

For investment, the model predicts less than half of the drop observed in the aftermath of the financial crisis. A key reason why agents do not slash investment is that the effective size of the capital stock was already diminished by the capital quality shock. The shock itself is like an enormous, exogenous disinvestment. Furthermore, our simple model abstracts from financial frictions which impeded investment (Gertler and Karadi (2011) and Brunnermeier and Sannikov (2014)).[22] A richer model with additional features and frictions will help bring the model closer to the data, but Figure 5 clearly demonstrates the quantitative potential of learning as a source of persistence, even in a standard business cycle setting.

Table 2 summarizes the long-run effects of the belief changes, by comparing the stochastic steady states associated with $\hat{g}_{2007}$ and $\hat{g}_{2009}$. It shows that capital and labor are 17% and 8% lower under the latter, which translates into a drop in output and consumption levels of about 12%. Investment is also lower by about 8%. Thus, even though the $\phi_t$ shocks experienced during the Great Recession were transitory, the resulting changes in beliefs permanently reduce economic activity.

**Credit spreads** The financial crisis had an large, immediate impact on credit spreads, but these effects were largely reversed within a few years. Some authors have argued that since heightened tail risk should inflate risk premia as well as credit spreads, the transitory effects of the financial crisis on credit spreads imply that tail risk plays at best a modest role in explaining macroeconomic stagnation (Hall, 2015). While the argument is intuitive, it ignores any endogenous response of investment or borrowing. A surge in risk triggers disinvestment and de-leveraging. Because firms borrow less, they default less, which offsets the increase in the credit spread. In our calibrated model, this offsetting force is very powerful. Table 2 reports the change in credit spreads, defined as the implied interest on risky debt, $1/q_t$ less the risk-free rate $r^f$. The credit spread in the stochastic steady state under the 2009 belief is only 4 basis points

---

[21]In Appendix B.3, we show that including shock realizations post-2009 does not materially change this finding.

[22]Alternative amplification mechanisms are studied in Adrian and Boyarchenko (2012); Jermann and Quadrini (2012); Khan et al. (2014); Zetlin-Jones and Shourideh (2014); Bigio (2015); Moriera and Savov (2015) among others.

higher than under the 2007 beliefs. Thus, belief revisions can have significant and long-lived real effects, even if the long-run change in credit spreads is very small.

| | Stochastic steady state levels | | Change |
|---|---|---|---|
| | $\hat{g}_{2007}$ | $\hat{g}_{2009}$ | |
| Output | 6.37 | 5.67 | -12 % |
| Capital | 27.52 | 22.80 | -17% |
| Investment | 0.71 | 0.66 | - 7% |
| Labor | 2.40 | 2.20 | -8% |
| Consumption | 5.66 | 5.01 | -12% |
| Credit spread (bps) | 84 | 88 | +4 |

Table 2: **Belief changes stemming from 2008-'09 shocks lead to significant reductions in economic activity.**

*Steady state numbers represent the average levels in the stochastic steady state associated with $\hat{g}_{2007}$ and $\hat{g}_{2009}$ respectively.*

**Capital and labor wedges**  Our results are also consistent with measurement strategies that impute 'wedges' or implicit taxes in the investment and labor optimality equations, in the spirit of the business cycle accounting framework laid out in Chari et al. (2007). These exercises, applied to the recent US experience, show persistently higher wedges in both capital and labor after the financial crisis. Hall (2015) documents a rise in the spread between return on capital and the risk-free rate, which he interprets as a rise in the capital wedge[23]. Karabarbounis (2014) also finds a significant increase in the implicit tax on labor. Our calibrated model predicts increases in both wedges from the increase in tail risk. The capital wedge, measured as in Hall (2015), in the new stochastic steady state (i.e. under $\hat{g}_{2009}$), is 1% higher than the 2007 one, while the implicit tax on labor, under the assumptions in Karabarbounis (2014) [24] goes up by 12 % (an increase of one-third relative to 2007-levels).

**Turning Off Belief Updating**  To demonstrate the role of learning, we plot average simulated outcomes from an otherwise identical economy where agents know the final distribution $\hat{g}_{2009}$ with certainty, from the very beginning (dashed line in Figure 5). These agents do not revise their beliefs. This corresponds to a standard rational expectations econometrics approach, where agents are assumed to know the true distribution of shocks hitting the economy and the econometrician estimates this distribution using all the available data. The post-2009 paths are

---

[23]Note that this includes a risk premium component as well as other (e.g. financial) factors.

[24]Karabarbounis (2014) uses a utility specification $u(C_t, L_t) = \frac{1}{1-\gamma}(C_t^{\eta}(T - L_t)^{1-\eta})^{1-\gamma}$ and a Cobb-Douglas production function and defines $\tau_t^{LW} = \log MPL_t - \log MRS_t$. He uses various data series to measure these objects and finds that the wedge increased by 25-50%.

simulated as follows: 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. After 2009, the sequence of shocks is drawn from the estimated 2009 distribution.

In the absence of belief revisions, the negative shock leads to an investment boom, as the economy seeks to replenish the lost effective capital. While the curvature in utility moderates the speed of this transition to an extent, the overall pattern of a steady recovery back to the original steady state is clear.[25] This shows that learning is what generates long-lived reductions in economic activity.

## 5.1 Shock Size and Persistence

The main question of the paper is not just whether belief revisions can help generate persistent effects from transitory shocks. It is also about why the effect of some recessions on macro variables is more long-lived than others. All belief changes are permanent. But this does not imply that all shocks have similar consequences. The total effect of an adverse shock to capital quality are a combination of the transitory direct effect and the permanent effect through beliefs. The extent to which a shock generates persistent outcomes depends on the relative size of these two effects.

Recall from Figure 4 that the adverse shocks observed during the 2008-'09 period were exceptionally large - the shock in 2009, for example, was almost 5 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.[26] The 1 standard deviation 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 1 standard deviation 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

---

[25] 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).

[26] For the remainder of the paper, we use an approximation which yields considerable computational tractability. Instead of simulating by drawing time paths and re-estimating beliefs, we compute impulse responses under the assumption that, after 2009, (i) beliefs remain fixed at $\hat{g}_{2009}$, (ii) leverage is constant and (iii) the shock realization at every date is equal to its average value, $E_t(\phi_t)$. We show in Appendix B.4 that the impulse responses under these restrictions are practically identical to the ones obtained through our full-blown simulation procedure. In light of this, we refer to these approximate impulse responses as our baseline in the remaining figures in the paper.

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 is not really about the size of the shock per se, but the effect it has on beliefs. If a large shocks had been observed frequently, it could also have transitory effects. More generally, our learning mechanism explains why fluctuations triggered by rare events (like the Great Recession) are particularly persistent.
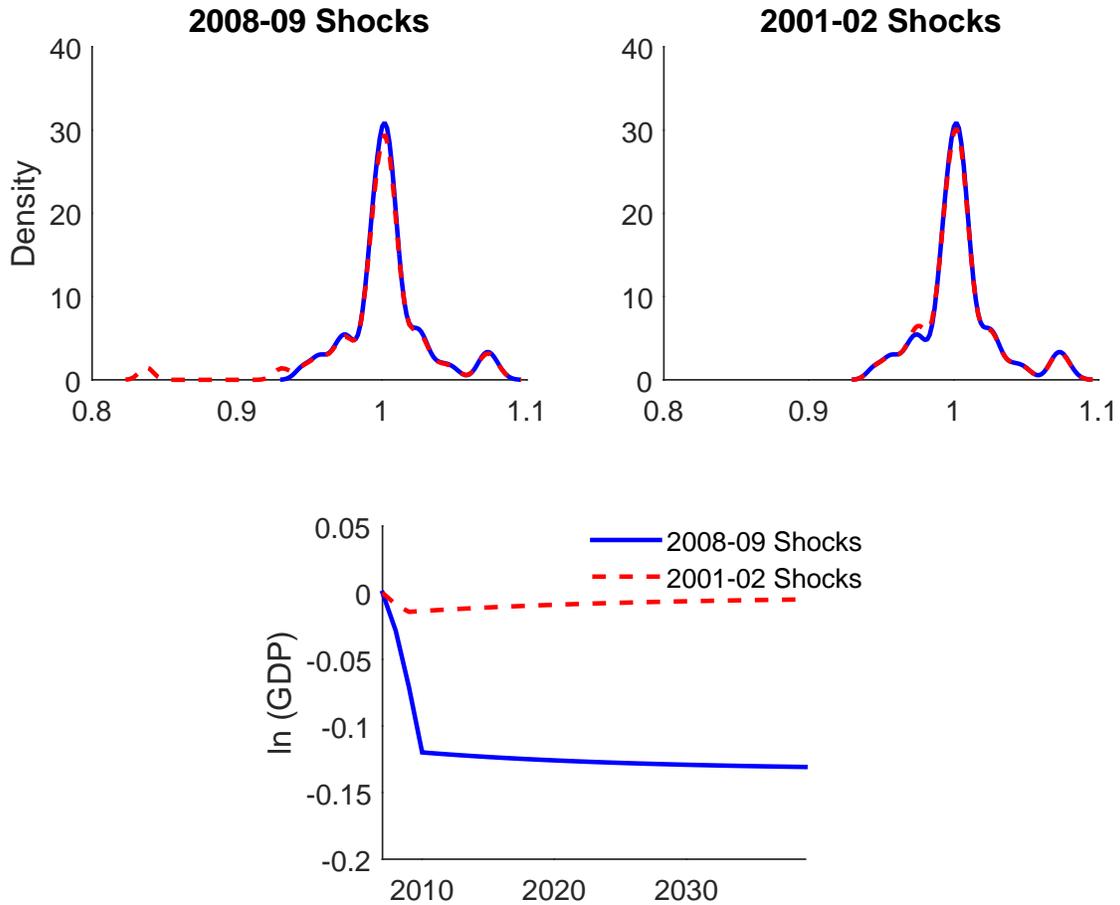


Figure 6: **Small shocks create negligible belief revisions.**
*The solid blue line in both top panels shows the estimated density in 2007, while the dashed lines show the new estimate after the 2008-09 shocks (left panel) and a counter-factual 2001-02 shocks (right panel). The bottom panel shows the response of output under the two different scenarios.*

## 5.2 Longer data sample and the Great Depression

Of course, the 2008-09 financial crisis was not the first time in history that the U.S. economy has experienced large adverse economic shocks. Since our simulations start in 1950, the Great Depression is not in our agents' information set. This raises an obvious question - how would

access to more data, with large adverse shocks, affect the response of beliefs to the recent financial crisis?

Once we include data from a different era with different institutions, the assumption that agents incorporate old and new data in their beliefs in the same way becomes less realistic. We consider the possibility that agents underweight past observations relative to more recent ones. This could be motivated by the possibility of unobserved regime shifts. We follow the learning literature in addressing this by allowing for the possibility of discounting.[27] This is also a common assumption in models of experiential learning with overlapping generations.

In the absence of discounting, i.e. if past data are treated on par with more recent observations, then, as the observed data series lengthens, beliefs converge, and eventually, new data ceases to affect beliefs. Our analysis in Section 2 suggests that beliefs about tail events are likely to converge more slowly than those elsewhere in the distribution, because of infrequent observations. Even so, convergence is inevitable. However, if agents underweight past data, beliefs will not converge. With approximated data extending back to the 19th century and modest data discounting, the belief changes induced by the 2008-09 experience continue to have a large, persistent effect on economic activity.

The most direct way to answer this question would be to give our agents a longer data sample which includes the 1930s. The problem is that the non-financial asset data that we use to construct $\phi_t$ is available only for the post-WW II period. Our first attempt to deal with the lack of data was to come up with an indirect measure. To this end, we projected the measured $\phi_t$ series post-1950 on a number of variables and used the estimated coefficients to impute values for $\phi_t$ pre-1950. However, this did not quite work, despite a fairly comprehensive search.[28] Our second, and admittedly cruder, attempt was to use the post-WW II sample to construct scenarios for what the pre-WW II sample might have looked like. Given that our goal is not to explain what happened during the Great Depression years, but rather, to understand how having more data affects learning today, this seems a reasonable place to start. We begin with the simplest exercise: Replicate the 1950-2009 sample. Specifically, assume that $\phi_t$ realizations for the period from 1890-1949 were identical to those in 1950-2009, with one adjustment: For 1929-30, we use multiples of the 2008-09 realizations. We set $\{\phi_{1929}, \phi_{1930}\} = \varepsilon \cdot \{\phi_{2008}, \phi_{2009}\}$, with $\varepsilon \geq 1$. In other words, we assume that the Great Depression saw adverse shocks that were scaled versions of those experienced during the Great Recession. We then repeat our analysis under the assumption that agents are endowed with this expanded data series (starting from

---

[27]See, for example, Sargent (2001), Cho et al. (2002) and Evans and Honkapohja (2001).

[28]We used a range of macro and asset pricing variables for this exercise - including GDP, unemployment, S&P returns and the Case-Shiller index of home prices. We also experimented with lead-lag structures to overcome the somewhat contemporaneous correlation with our measured $\phi_t$ series. Across specifications, the resulting projections for 1929-1930 showed only modestly adverse realizations.

1890). Now, when the financial crisis hits, the effect on beliefs is moderated by two factors - one, the larger data sample, which reduces the influence of new observations and two, the fact that agents had 'seen' a similar episode before.

To capture discounting, we modify our kernel estimation procedure to allow for weights. Observation from $s$ periods earlier are assigned a weight $\lambda^s$, where $\lambda \leq 1$ is a parameter. When this is set to 1, there is no discounting of past observations.[29]

The first panel of Figure 7 shows the impulse response without discounting (i.e. with $\lambda = 1$). It shows a slow increase in output following the crisis and some attenuation of the long-run effect, but the changes relative to the baseline case in Figure 5 are modest. When older data is discounted by 1% ($\lambda = 0.99$, the center panel), this attenuation almost completely disappears and the impulse responses lie right on top of our baseline estimates.[30]

Perhaps the true magnitude of the Great Depression shocks is far larger than those seen in 2008-09. To explore this possibility, we redo the analysis with shocks in 1929-30 that are twice as bad as those in 2008-09. This implies that $(\phi_{1929}, \phi_{1930}) = (0.86, 0.70)$. Note that these are very large shocks - they are respectively 5 and 10 standard deviations below the mean and taken together, imply an erosion of almost 50% in the stock of effective capital for the entire US economy. The results from this version (with the data discounting parameter $\lambda = 0.99$) are presented in the third panel of Figure 7. Again, the long-run effects of the recession are attenuated, but only modestly.

In sum, these results suggest that expanding the information set by adding more data does not drastically alter our main conclusions, especially once we allow for the possibility of discounting older data.

# 6 Understanding the Economic Response to Belief Changes

How does the model generate such large fluctuations in response to changes in probabilities that are still quite small? To answer this, we turn off features of the model one-by-one in order to isolate how much each one contributes to the results in Figure 5. First, we analyze the contribution of changes in the first moment of the distribution (i.e. in the average $\phi$). Next,

---

[29]When $\lambda < 1$, eventually, the 2008-09 observations will be discounted as well. However, since we are always drawing from $\hat{G}_{2009}$, averaging across sample paths still yields $\hat{G}_{2009}$. Since our primary goal here is to show the effect of additional data and the spirit of our exercise is to put economist and agent on the same footing, this was a natural assumption.

[30]If we keep increasing the discount, the decline in long-run GDP becomes *bigger*. Intuitively, with high enough discounting, the weight of recent observations increases beyond the level in the undiscounted, reduced sample used for our baseline analysis. For example, with $\lambda = 0.98$, GDP drops by about 16% in the new steady state.
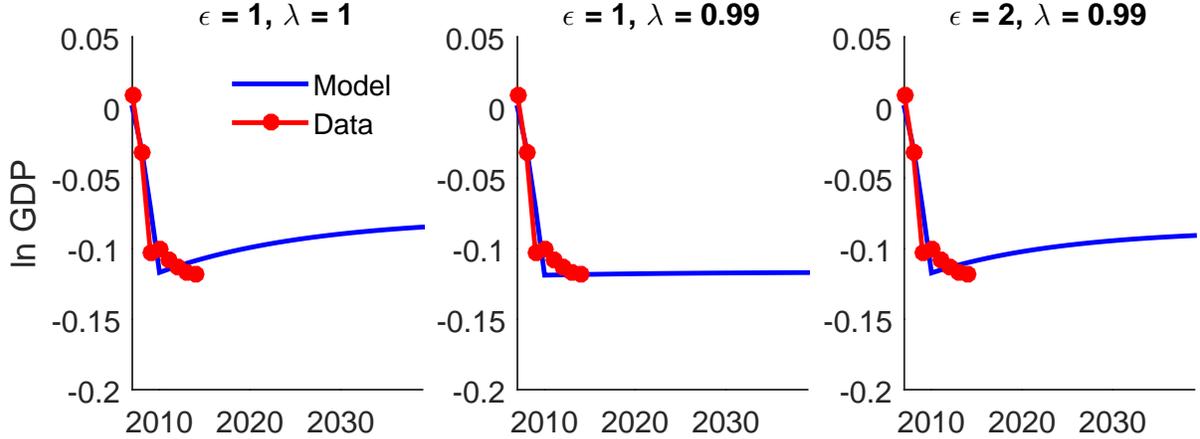
Figure 7: **Extending the data sample does not materially change the effect of belief revisions.**

*Each panel plots the response of GDP to the 2008-09 shocks under a hypothetical information set, starting from 1890. To fill in the data for the period 1890-1949, we use the observed time series from 1950-2009, with $\{\phi_{1929}, \phi_{1930}\} = \varepsilon \cdot \{\phi_{2008}, \phi_{2009}\}$. The parameter $\lambda$ indexes the extent to which older observations are discounted where $\lambda = 1$ represents no discounting.*

we quantify the role of risk aversion by analyzing a version with quasi-linear preferences. The third experiment shows the interaction between learning and debt by comparing our results to an economy with no leverage. The bottom panel of Table 3 summarizes the results of these counterfactual exercises. It shows that all 3 components - learning, risk aversion and debt play a significant role in generating persistent stagnation. Removing any of these elements would eliminate between one-fourth and one-half of our long-run effects.

## 6.1 Role of changes in mean $\phi$

We decompose the total effect of belief revisions into a component attributable to changes in the mean (or average $\phi$) and the remaining attributable to changes in higher moments. To do this, we adjust the estimated distribution in 2009 so that $\mathbb{E}_{2009}(\phi_t) = \mathbb{E}_{2007}(\phi_t)$. The dashed line in the first panel of Figure 8 shows the response of output under this specification. Even with the mean change taken out, the long-run fall in GDP is about 6%, about half of the total effect in our baseline case (the solid lines). In other words, while the effects coming from mean changes are significant, the contribution of higher moment changes is just as important in generating the persistent decline in economic activity from belief revisions.

The change in the mean $\mathbb{E}_t[\phi_t]$ between 2007 and 2009 is relatively modest, only about 0.4%, but, as the graph shows, its effect on long-run is about 6%. To dig a little deeper into why long-run outcomes are so sensitive to $\phi$, we turn to a special case - a deterministic version of

|  | 2014 | Long run |
|---|---|---|
| Data | −0.12 | |
| Benchmark model | −0.12 | −0.12 |
| *Counterfactuals: Beliefs* | | |
| No learning | −0.09 | 0.00 |
| Small shock | −0.02 | 0.00 |
| Longer data sample ($\epsilon = 1$, $\lambda = 1$) | −0.12 | −0.09 |
| Longer data sample ($\epsilon = 1$, $\lambda = 0.99$) | −0.12 | −0.12 |
| Longer data sample ($\epsilon = 2$, $\lambda = 0.99$) | −0.12 | −0.09 |
| *Counterfactuals: The economic model* | | |
| Constant mean | −0.09 | −0.06 |
| No risk aversion | −0.07 | −0.07 |
| No debt | −0.12 | −0.09 |

Table 3: **Change in GDP relative to 2007 steady state.**

our economy without debt. The level of steady state capital is given by the following equation[31]

$$\ln k_{ss} = \text{Const.} + \left( \frac{1+\gamma}{\gamma} \frac{\alpha}{1-\alpha} \right) \ln \phi_{ss} - \left( \frac{1}{1-\alpha} \frac{\alpha+\gamma}{\gamma} \right) \ln \left( \frac{1}{\beta} - (1-\delta) \phi_{ss} \right). \qquad (16)$$

Hence, the effect of the mean shock on steady sate capital is given by

$$\frac{d \ln k_{ss}}{d \ln \phi_{ss}} = \left( \frac{1+\gamma}{\gamma} \frac{\alpha}{1-\alpha} \right) + \left( \frac{1}{1-\alpha} \frac{\alpha+\gamma}{\gamma} \right) \frac{(1-\delta)}{1/\beta - (1-\delta)\phi_{ss}}.$$

Under our parameterization,

$$\frac{1+\gamma}{\gamma} \frac{\alpha}{1-\alpha} = 2, \qquad \frac{1}{1-\alpha} \frac{\alpha+\gamma}{\gamma} = 3, \qquad \left( \frac{(1-\delta)}{1/\beta - (1-\delta)\phi_{ss}} \right)_{\phi_{ss}=1} = 7.5$$

which implies $\frac{d \ln k_{ss}}{d \ln \phi_{ss}} = 2 + 3(7.5) = 24.5$. This simple calculation shows the source of the high sensitivity - the fact that capital quality shock affects not just the current return component but also the portion that comes from the undepreciated stock. In other words, it has a bigger effect on the total return to capital than say, total factor productivity. This combined with the high elasticity of steady state capital lies behind the significant effects of the change in the mean $\phi$.

---

[31]In steady state, $M_t = 1$ and the intertemporal Euler equation and labor optimality conditions reduce to

$$\begin{aligned} 1 &= \beta \left( \alpha \phi_{ss}^\alpha k_{ss}^{\alpha-1} l_{ss}^{1-\alpha} + \phi_{ss} (1-\delta) \right) \\ l_{ss}^\gamma &= W_{ss} = (1-\alpha) \phi_{ss}^\alpha k_{ss}^\alpha l_{ss}^{-\alpha}. \end{aligned}$$

Substituting for $l_{ss}$ from the second into the first and re-arranging yields the expression (16).
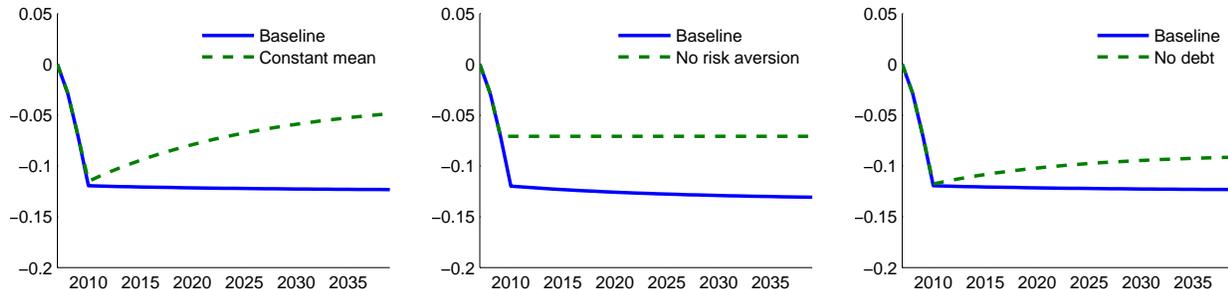
Figure 8: **Understanding the role of expected return, risk aversion and debt.**
*Change in ln GDP under 3 different scenarios: (1) the mean capital quality shock $E[\phi_t]$ is held fixed; (2) no risk aversion, and (3) no debt.*

## 6.2   Role of risk aversion

Next, 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. This exercise allows us to see how much of the persistent drop in investment and hiring comes from the interaction of curvature in utility and changes in beliefs about returns to investing.

The second pannel of Figure 8 presents the time path for output in this version. 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). The graph indicates that this effect is quite strong and accounts for almost half of the long-run drop in output.

## 6.3   Role of debt

Finally, we examine the interaction between learning and debt by comparing our results to an identical economy where all investment is financed through equity. Formally, we set the tax advantage parameter $\chi$ to 1 and so leverage is equal 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.

The third pannel of Figure 8 plots the time path for output for this variant of our model, along with our baseline version from Figure 5. The graph shows that the effects of belief revisions are smaller in the absence of debt though still significant - they lead to a 9% long-run reduction in output, compared to 12% in the version with debt. Thus, the presence of

defaultable debt amplifies the effects of changes in tail risk and contributes about a fourth of the long run macroeconomic response.

**The Interaction of Debt and Large vs Small Shocks**  Debt also plays an important role in one of the main questions of the paper, namely why some shocks generate more persistent responses than others. The attractiveness of debt (and therefore, the incentives to borrow) is affected disproportionately by perceived tail risk - and since larger shocks changes belief further out in the tail, they are amplified to a much greater extent by debt. In contrast, smaller shocks induce belief changes closer to the mean and therefore, do not interact with debt in a significant fashion. This is demonstrated in Figure 9.
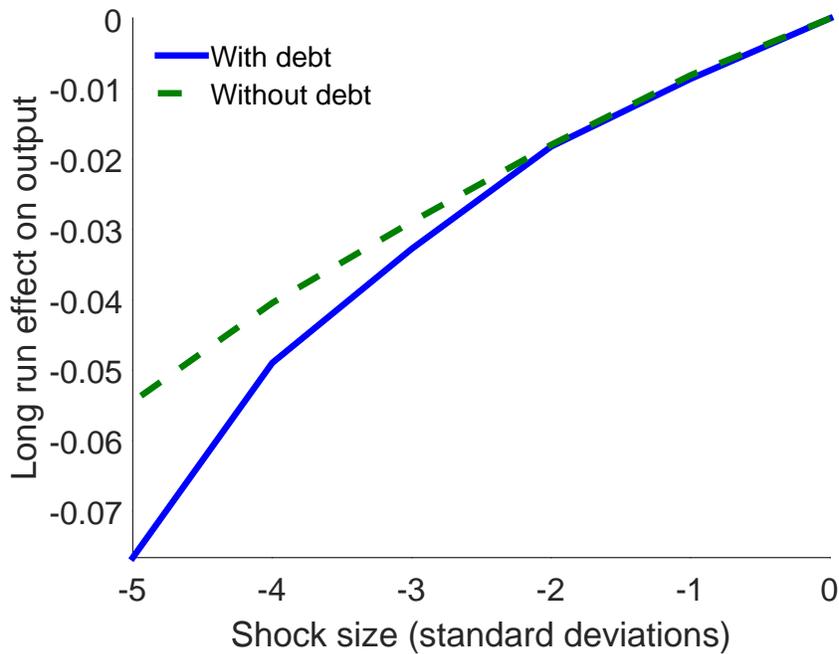


Figure 9: **Debt amplifies belief revisions from large shocks.**
*The graph shows the change in long run GDP both with (solid line) and without debt (dashed line) in response to negative shocks of various sizes. The initial condition is the steady state associated with $\hat{g}_{2007}$.*

We subjected our model economy calibrated to the 2007 beliefs with shocks ranging in size from 1 to 5 standard deviations. Figure 9 plots the corresponding effects on GDP in the long-run (i.e. at the new stochastic steady state). It shows that debt adds significant non-linearity to the response of the economy to large shocks - the responsiveness to small shocks is pretty much almost the same as without debt, but larger shocks see significant amplification.

To reiterate, debt is not necessary for transitory shocks to have some persistent effect. Belief updating alone will do that. What debt does is make tail risk revisions look very different from other belief changes. This contributes to making the long-run effects from extreme recessions

very different from responses to milder downturns.

# 7    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

Adrian, T. and N. Boyarchenko (2012). Intermediary leverage cycles and financial stability. (567).

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

Anzoategui, D., D. Comin, M. Gertler, and J. Martinez (2015). Endogenous technology adoption and r & d as sources of business cycle persistence. New York University working paper.

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.

Bigio, S. (2015). Endogenous liquidity and the business cycle. *American Economic Review.*

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).

Chari, V., P. J. Kehoe, and E. R. McGrattan (2007). Business cycle accounting. *Econometrica 75*(3), 781–836.

Cho, I.-K., N. Williams, and T. J. Sargent (2002). Escaping nash inflation. *The Review of Economic Studies 69*(1), 1–40.

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

Cooley, T. and E. Prescott (1995). *Frontiers of Business Cycle Research*, Chapter Economic Growth and Business Cycles. Princeton University Press.

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.

Evans, G. W. and S. Honkapohja (2001). *Learning and expectations in macroeconomics*. Princeton University Press.

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.

Gomme, P., B. Ravikumar, and P. Rupert (2011). The return to capital and the business cycle. *Review of Economic Dynamics 14*(2), 262–278.

Gordon, R. (2014). The demise of u. s. economic growth: Restatement, rebuttal, and reflections. NBER working paper 19895.

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.

Hall, R. (2015). Macroeconomics of persistent slumps. *Handbook of Macroeconomics*.

Hansen, B. E. (2015). *Econometrics*.

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.

Jermann, U. and V. Quadrini (2012). Macroeconomic effects of financial shocks. *American Economic Review 102*(1), 238–71.

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

Karabarbounis, L. (2014). The labor wedge: Mrs vs. mpn. *Review of Economic Dynamics 17*(2), 206–223.

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.

King, R., C. Plosser, and S. Rebelo (1988). Production growth and business cycles. *Journal of Monetary Economics 21*, 195–232.

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

Maćkowiak, B. and M. Wiederholt (2010). Business cycle dynamics under rational inattention. *Review of Economic Studies 82*(4), 1502–1532.

Merton, R. (1973). An intertemporal capital asset pricing model. *Econometrica 41*(5), 867–887.

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.

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

Sargent, T. J. (2001). *The conquest of American inflation.* Princeton University Press.

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.

Zetlin-Jones, A. and A. Shourideh (2014). External financing and the role of financial frictions over the business cycle: Measurement and theory. Working paper.

# Appendix

## A    Optimality conditions from firm's problem

The problem of the firm is

$$V\left(\Pi_{it}, B_{it}, S_t\right) = \max\left[0, \max_{d_{it}, \hat{k}_{it+1}, b_{it+1}, w_{it+1}, l_{it+1}} d_{it} + \mathbb{E}_t M_{t+1} V\left(\Pi_{it+1}, B_{it+1}, S_{t+1}\right)\right] \qquad (17)$$

subject to

$$\begin{aligned}
\text{Dividends:} \quad & d_{it} \leq \Pi_{it} - B_{it} - \hat{k}_{it+1} + \chi q_{it} b_{it+1} \\
\text{Discounted wages:} \quad & \mathcal{W}_t \leq w_{it+1} q\left(\hat{k}_{it+1}, l_{it+1}, B_{it+1}, S_t\right) \\
\text{Future obligations:} \quad & B_{it+1} = b_{it+1} + w_{it+1} l_{it+1} \\
\text{Resources:} \quad & \Pi_{it+1} = v_{it+1}\left[A(\phi_{t+1}\hat{k}_{it+1})^\alpha l_{it+1}^{1-\alpha} + (1-\delta)\phi_{t+1}\hat{k}_{it+1}\right] \\
\text{Bond price:} \quad & q\left(\hat{k}_{it+1}, l_{it+1}, B_{it+1}, S_t\right) = \mathbb{E}_t M_{t+1}\left[r_{it+1} + (1-r_{it+1})\frac{\theta \tilde{V}_{it+1}}{B_{it+1}}\right] \\
& r_{it+1} = \begin{cases} 0 & \text{if } v_{it} < \underline{v}(S_t) \\ 1 & \text{if } v_{it} \geq \underline{v}(S_t) \end{cases}.
\end{aligned}$$

We begin by noting that the firm's problem can be expressed in terms of choosing capital, leverage, defined as $lev_{it+1} \equiv \frac{B_{it+1}}{\hat{k}_{it+1}}$, and the labor-capital ratio, $\frac{l_{it+1}}{\hat{k}_{it+1}}$. Define

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

Then, the firm's maximization problem becomes

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

and

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

We guess and then verify that $\Gamma_{it} = 0$.[32] Replacing the debt price schedule and rearranging terms yields

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

$$\tilde{J}_{t+1}^k = R_{t+1}^k + lev_{it+1} (\chi - 1) r_{t+1} + (\chi\theta - 1)(1 - r_{t+1}) v_{it} R_{t+1}^k.$$

By definition of the default threshold, we have $\mathbb{E} r_{t+1} = (1 - F(\underline{v}))$. Also, note that the default threshold becomes $\underline{v} = \frac{lev_{it+1}}{R_{t+1}^k}$. Hence

$$\tilde{J}_{t+1}^k = R_{t+1}^k \left( 1 + \underline{v} (\chi - 1)(1 - F(\underline{v})) + (\chi\theta - 1) h(\underline{v}) \right)$$

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

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

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

$$\underline{v} = \frac{lev_{it+1}}{R_{t+1}^k}$$

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

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

which implies equation (12) 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

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

---

[32]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

$$R_{t+1}^k \frac{\partial J^k(\underline{v})}{\partial \frac{l_{t+1}}{\hat{k}_{it+1}}} = R_{t+1}^k \frac{\partial \underline{v}}{\partial \frac{l_{t+1}}{\hat{k}_{it+1}}} \left( (\chi - 1)(1 - F(\underline{v})) - \underline{v}(\chi - 1) f(\underline{v}) + (\chi \theta - 1) \frac{\partial h(\underline{v})}{\partial \underline{v}} \right)$$

$$\frac{\partial \underline{v}}{\partial \frac{l_{t+1}}{\hat{k}_{it+1}}} = -\frac{lev_{it+1}}{(R^k)^2} \frac{\partial R^k}{\partial \frac{l_{t+1}}{\hat{k}_{it+1}}} = -\frac{\underline{v}^2}{lev_{it+1}} \frac{\partial R^k}{\partial \frac{l_{t+1}}{\hat{k}_{it+1}}}$$

$$\frac{dh(\underline{v})}{d\underline{v}} = \underline{v} f(\underline{v})$$

$$\frac{\partial R_{t+1}^k}{\partial \frac{l_{t+1}}{\hat{k}_{it+1}}} = v_{it} A (1 - \alpha) \phi_{t+1}{}^\alpha \left( \frac{l_{it+1}}{\hat{k}_{it+1}} \right)^{-\alpha}.$$

Rearranging terms yields

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

$$J^l(\underline{v}) = 1 + \underline{v}^2 f(\underline{v}) \chi (1 - \theta) - (1 - \chi \theta) h(\underline{v}),$$

which is (13) in the main text.

Finally, the first order condition with respect to leverage is

$$\mathbb{E} M_{t+1} R_{t+1}^k \frac{\partial J_{t+1}^k}{\partial lev_{it+1}} = 0,$$

where

$$\frac{\partial J_{t+1}^k}{\partial lev_{it+1}} = \frac{\partial \underline{v}}{\partial lev_{it+1}} \left( (\chi - 1)(1 - F(\underline{v})) - (\chi - 1) \underline{v} f(\underline{v}) + (\chi \theta - 1) \underline{v} f(\underline{v}) \right)$$

$$= \frac{1}{R_{t+1}^k} \left( (\chi - 1)(1 - F(\underline{v})) - \chi (1 - \theta) \underline{v} f(\underline{v}) \right)$$

hence

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

which is (14) in the main text.

# B    Additional Results

## B.1    Measurement of $\phi_t$: Alternative price indices

Figure 10 shows that the measurement of capital quality shocks is unaffected when we use the price index for GDP or Personal Consumption Expenditure to control for nominal price changes.

## B.2    Numerical accuracy of solution method

To test the numerical accuracy of our solution method we perform the following exercise. Starting from the steady state of $\hat{g}_{2007}$, we simulate time paths for two different economies. In Model I, as new data arrives, we
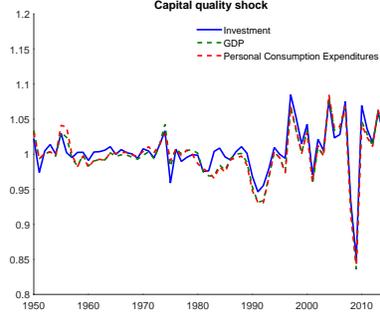
**Capital quality shock**

Figure 10: **Time series of measured capital quality shocks using different indices to control for nominal price changes.**

update beliefs and policy functions at each date and history. In Model II, beliefs and policy functions are fixed at $\hat{g}_{2007}$. Our solution method essentially assumes that agents use Model II as an approximation for Model I, while evaluating continuation values. Table 4 shows the sample mean and coefficient of variation for output at different horizons for these two versions.[33] It is easy to see that aggregates (or at least, the first two moments thereof) are very well-approximated by replacing the sequence of future distributions with their conditional mean. Recall that this numerical procedure works reasonable well thanks to the martingale property of beliefs.

|  | Horizon | | | |
|---|---|---|---|---|
|  | $s = 1$ | $s = 5$ | $s = 10$ | $s = 15$ |
| $\mathbb{E}_t\left[y_{t+s}\right]$ | | | | |
| Model I: | 6.378 | 6.385 | 6.393 | 6.397 |
| Model II: | 6.378 | 6.385 | 6.394 | 6.398 |
| $\mathbb{CV}_t\left[y_{t+s}\right]$ | | | | |
| Model I: | 0.010 | 0.032 | 0.042 | 0.046 |
| Model II: | 0.010 | 0.031 | 0.040 | 0.044 |

Table 4: **Numerical accuarcy.**
*The rows labeled Model I show the actual moments under the assumption that beliefs $\hat{g}_{2007+s}$ are re-estimated at each date. Model II corresponds to the assumption underlying our solution method, where future beliefs are replaced by $\hat{g}_{2007}$.*

## B.3  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 estimated distribution at that date. For the period after 2014, we adopt the same approximation strategy, 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 11, along with the de-trended data. We note that the patterns implied by the model are quite close to the observed ones.

---

[33]These are averages over 4000 paths. Other aggregate variables, e.g. capital and labor, show similar patterns.
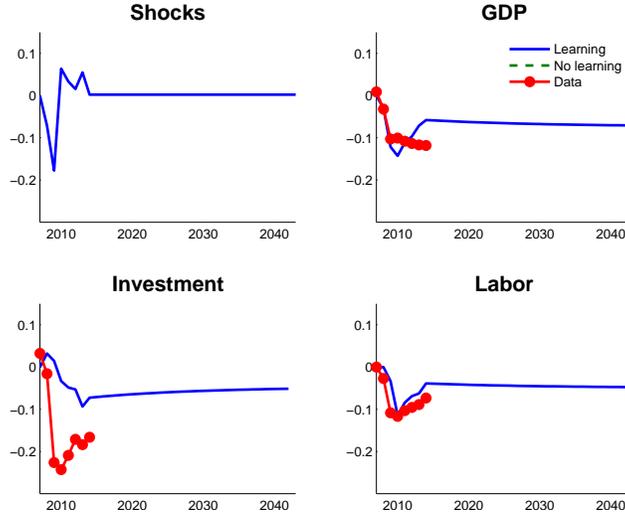
Figure 11: **Model vs data from 2008-2014.**
*Solid line is the baseline model subjected to the observed sequence of shocks from 2008-2014. The red circles are US data, in deviations from their pre-crisis trends.*

## B.4 Simulation vs Approximation

Recall that our baseline results in Figure 5 were computed by averaging across a large number of sample paths. Along each path, based on the draw of shocks, we re-estimated beliefs and derived corresponding policy functions. Here, we show that we can obtain a very good approximation of the impulse responses using a much simpler computational strategy. Specifically, for post-2009 dates, we hold fixed shock realizations (at their average value), beliefs (at $\hat{g}_{2009}$) and leverage. Note that this involves computing a single path using the 2009 policy functions (since beliefs are assumed to stay fixed at $\hat{G}_{2009}$). Figure 12 compares the resulting impulse response functions to those that emerge from the full blown simulation in our baseline. It shows that both strategies yield very similar paths for output and labor. For investment, our approximation slightly overstates the change, but the pattern is quite similar.
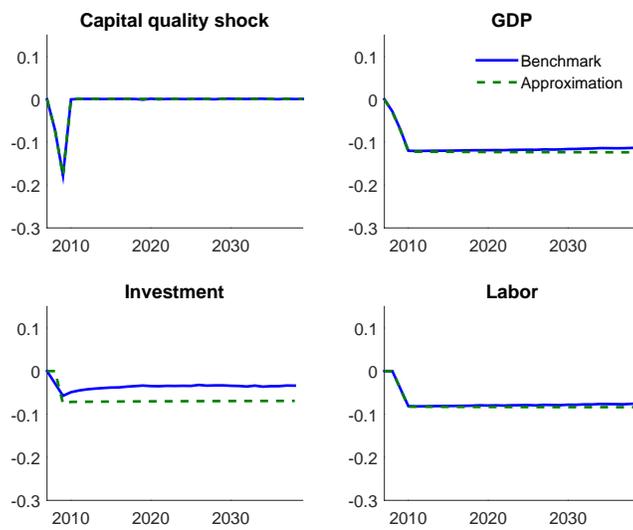
Figure 12: **Belief revisions post-2009.**

*Solid blue line shows baseline results computed by average across sample paths, with beliefs updated after draw. The dashed line is the approximate version with fixed leverage, beliefs and shock realizations.*