The Tail that Keeps the Riskless Rate Low

Julian Kozlowski
NYU

Laura Veldkamp
NYU Stern, CEPR, NBER

Venky Venkateswaran*
NYU Stern, NBER

August 14, 2017

Abstract

Riskless interest rates fell in the wake of the financial crisis and have remained low. We explore a simple explanation: This recession 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, perceived tail risk remains high, generates a demand for riskless, liquid assets, and continues to depress the riskless rate. To explore this idea, we study a simple production economy with liquidity constraints. Our twist on the model is that agents 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, the shock stays 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 drop in the risk-free rate.

JEL Classifications: D84, E32

Keywords: Interest rate, tail risk, belief-driven business cycles

*kozjuli@nyu.edu; lveldkam@stern.nyu.edu; vvenkate@stern.nyu.edu. We thank ... for their insightful discussions.

1 Introduction

Interest rates on riskless government debt have been falling in the aftermath of the financial crisis. While labor, production and consumption have rebounded, the riskless rate has remained stubbornly low. As Figure 1 shows, long-term interest rates fell following the financial crisis and never rebounded. We explore a simple explanation: the 2008-09 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. 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). These changes in beliefs about what kinds of tail events are possible endure long after the event itself has passed. This persistent increase in perceived tail risk makes liquid assets more valuable and thus continues to depress the yield on liquid government bonds today.

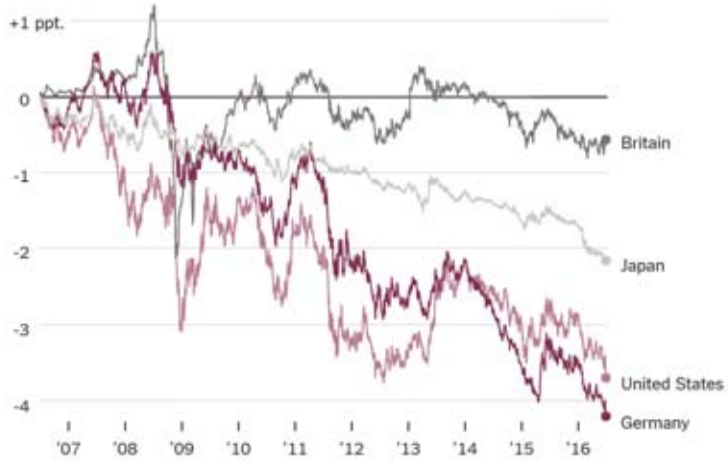


Figure 1: **Interest Rates Fell Following the Financial Crisis.** Change in percentage points of 10 -year government bond yield since July 3, 2006. A similar pattern emerges, even if we control for inflation. 10 -year rates typically reflect current fundamentals more than transitory monetary policy. Source: NYT June 28, 2016.

To measure how much tail risk rose, explain why it remained elevated, and explore its economic consequences, we incorporate our belief formation mechanism into a standard production economy with liquidity constraints. The non-standard element in the model is that instead of knowing the distribution of shocks, our agents use standard econometric tools to estimate the distribution, 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. The liquidity constraint in the model serves to amplify this effect

by making the yield on riskless, liquid assets more sensitive to the probability of a disaster. We feed a historical time-series of 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 post-crisis decline in the riskless rate.

The reason that heightened tail risk depresses the interest rate, or yield, on safe assets in our model is two-fold. First, an increase in risk alone makes riskless assets more valuable. Faced with more risk, agents want to save more. But not every agent can save more. The bond market has to clear. Therefore, the return on bonds declines in order to clear that market. This force alone explains about a third of the decline in the interest rate. The second force at work in our model is that safe assets offer liquidity services that are particularly valuable in very bad states. When the probability of these tail events rises, liquid assets are more valuable, and their yield declines to clear the market.

For most macro models, including belief-driven ones, producing persistent deviations of interest rates from trend requires sufficiently persistent shocks.[1] With an exogenous change in persistence, such a theory cannot explain *why* the rate has persistently remained low when fluctuations in previous cycles were more transitory. We show that when agents repeatedly re-estimate the distribution of shocks, persistence is endogenous and state-dependent. 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 and more permanent interest rate responses as well.

This view is quantitatively successful in explaining the observed patterns in interest rates and reconciling it with options data. 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

[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.

the end of World War II. (Huffington Post Oct.6, 2013)

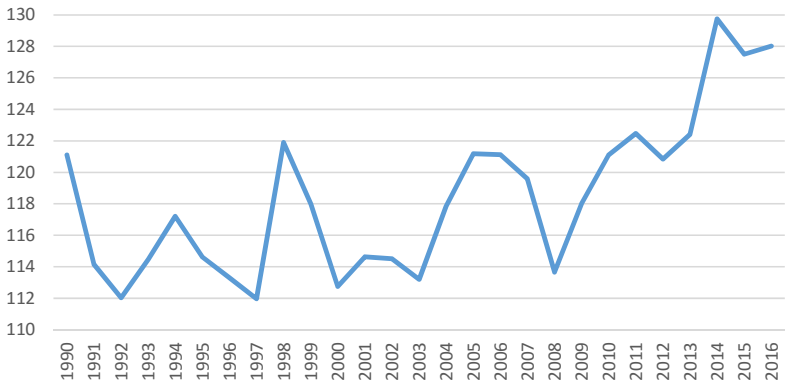


Figure 2: **The SKEW Index.**
A measure of the market price of tail risk on the S&P 500, constructed using option prices. Source: Chicago Board Options Exchange (CBOE). 1990:2016.

The decline in interest rates following the financial crisis took place in the context of a real rate that had been declining since the early 1980's. Obviously, the cause of this longer-run trend does not lie in the shock realizations of the financial crisis. Instead, it may come from a gradual change in expectations following the high inflation in the 1970's, or a surge in savings from emerging markets seeking safe assets to stabilize their exchange rates. This long-run phenomenon taking place in the background is hugely important, but distinct from our question. We seek to explain why, even relative to this downward trend, interest rates fell below trend in the financial crisis and failed to rebound.

We begin in Section 2 by examining the belief-formation mechanism in a simple representative agent problem, without liquidity constraints. 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 about the riskiness of the economy. Greater perceived risk makes riskless assets more valuable and depresses the riskless rate. 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 get a more realistic sense of the effect of tail risk on the riskless rate, Section 5 explores a model where riskless assets also offer liquidity services to facilitate real investment. The setting

is a standard production economy with a liquidity friction, similar to that in Ricardo Lagos (2005). Every period, there is a probability that any given agent has the option to produce a special good. When such production is possible, the agent wants to exploit the opportunity but requires liquid resources to do so. Riskless assets can always be used for special production. But for risky claims on the existing capital (equity), only a fraction can liquidated to use for special production in any given period. These liquidity constraints do not change the basic belief updating mechanism. Tail events have persistent effects for the same reason as in the simpler model. But the liquidity constraints serve to amplify the effect of a change in tail risk on the riskless rate.

Section 4 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 interest rates. Right now, this effect is small. Our work on progress is exploring how to amplify the size of the drop with liquidity constraints.

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, interest rates would fall and quickly rebound. 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.

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 **??** 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.

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 to introduce persistent effects of rarely observed shocks.

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.

Kozlowski et al. (2015) use a similar belief formation mechanism to explain persistence in real output fluctations. But it does not have any role for liquid assets and can therefore not explain the large decline in the riskless rate. Our belief formation process also has similarities 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 potential 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.

Finally, our paper contributes to the recent literature on low interest rates. For example, Bigio (2015) argues that a surge in asymmetric information about credit quality can explain the fall in the riskless rate. But the papers in this literature do not explain the persistence of this low rate. Our model offers a mechanism to produce persistent effects from transitory tail events.

[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).

2 A Representative Agent Economy

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 about tail risk, why they are persistent, and how this reduces the riskless rate. The mechanism 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.

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

Belief Formation 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 ϕ_t whose true density g is unknown to agents in the economy. The agents do know that the shock ϕ_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 ϕ_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, κ_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 ϕ_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 agent expects her future belief to be the same as her current beliefs. This property holds exactly

[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).

if the bandwidth parameter κ_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.

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 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.

3.1 Measuring capital quality shocks

To construct a time series of ϕ_t, we use two series on non-financial assets in the US economy from the Flow of Funds reports. One series is 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

[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 ϕ, $j \geq 1$

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

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

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 $\mathbb{E}_t\left[\hat{G}_{t+j}^0\left(\phi\right)\middle|\mathcal{I}_t\right] = \hat{G}_t^0\left(\phi\right)$.

[6] See Hansen (2015).

their model counterparts as follows

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

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 $PINDX_t$).[7] This allows us to recover the quality shock ϕ_t

$$\begin{aligned} \phi_t = \frac{K_t}{\hat{K}_t} &= \left(\frac{P_t^k K_t}{P_{t-1}^k \hat{K}_t} \right) \left(\frac{P_{t-1}^k}{P_t^k} \right) \\ &= \left(\frac{NFA_t^{RC}}{(1-\delta) NFA_{t-1}^{RC} + NFA_t^{HC} - (1-\delta)\, NFA_{t-1}^{HC}} \right) \left(\frac{PINDX_{t-1}^k}{PINDX_t^k} \right) \end{aligned} \tag{1}$$

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 (1), we construct an annual time series for capital quality shocks for the US economy since 1950. The left panel of Figure 3 plots the resulting series. For most of the sample period, the shock realizations are in a relatively tight range around 1, but at the onset of the recent Great Recession, we saw two large adverse realizations: 0.93 in 2008 and 0.84 in 2009. To put 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 3. 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.

3.2 Calibration

A period is interpreted as a year. We choose the discount factor $\beta = 0.95$ and depreciation $\delta = 0.06$ (consistent with the reported depreciation rates in the Flow of Funds data) to target a steady state capital-output ratio of 3.5 (this is taken from Cooley and Prescott (1995)). The

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

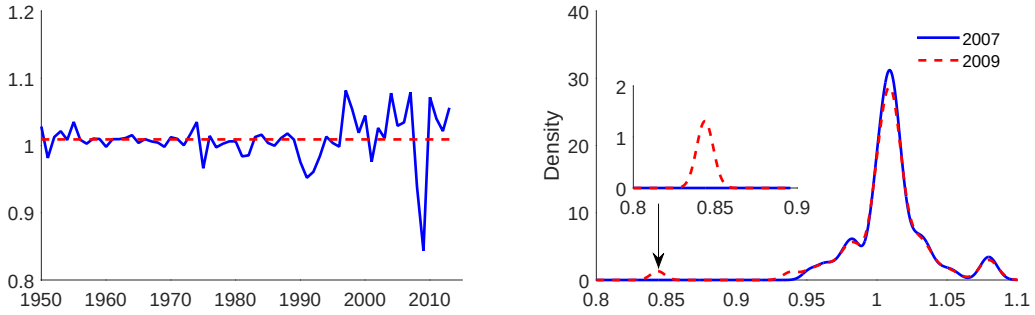


Figure 3: **Capital quality shocks.**

The left panel shows the time series of ϕ_t measured from the US data using (1). 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.

share of capital in the production, is 0.40, which is also taken from Cooley and Prescott (1995). The labor supply parameter, γ, is set to 0.5, in line with Midrigan and Philippon (2011), corresponding to a Frisch elasticity of 2. The labor disutility parameter ζ and the TFP term in production are normalized to 1. Finally, the risk aversion parameter is set to 2. Table 1 summarizes the resulting parameter choices.

Parameter	Value	Description
Preferences:		
β	0.95	Discount factor
σ	2	Risk aversion
γ	0.50	1/Frisch elasticity
ζ	1	Labor disutility
Technology:		
α	0.40	Capital share
δ	0.06	Depreciation rate
A	1	TFP

Table 1: **Parameters**

4 Results

Our main goal in this section is to quantify the size and persistence of the riskless interest rate response to a large but transitory shock ϕ_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 ϕ_t from 1950-2007, measured using the strategy outlined in Section 3 and

computing the stochastic steady state associated with this distribution.[8] 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.

4.1 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.

Note that the left tail "hump" never disappears if future shocks are drawn 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 (ϕ_t), even ones that are transitory, has a permanent expected effect on beliefs. Sometimes the permanent effects are small and most of the effect is transitory. But sometimes, when the shock is one that has rarely or never been observed before, the permanent effect is large.

4.2 Representative Agent Results

With this model, we perform the following experiment. We feed in the capital quality data from 1950-2007 and allow agents to form their beliefs. Given these beliefs, we start the model is its steady state. Then, we hit the economy with the capital quality shocks from 2008 and 2009. We re-compute the steady state. The results in Table 2 compare the aggregates in the 2007 steady state to the 2009 steady state.[9]

[8] 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.

[9] We are working on a more sophisticated dynamic exercise: After 2009, we want to observe the long run evolution of the economy, an in particular its interest rate. But we only have a few more years of data. To observe long run behavior of the model, we draw random 30-year sequences of capital quality shocks from the 2009 estimated distribution. For each sequence, we solve the model at each date. Then, we average the macro aggregates from all the simulations, on each date. That average response is what we plan to report.

	$\hat{g}_{2007}$	$\hat{g}_{2009}$	Change
Interest Rate	5.27	5.12	15 bps
Output	6.078	5.768	5.12%

Table 2: Steady State Interest Rates and Macro Aggregates, Pre- and Post-Crisis

Results from representative agent model (no liquidity constraints).

5 Model with Liquidity Constraints

The model analyzed in the previous sections highlights the ability of our belief revision mechanism to induce long-lived changes in the real interest rate. However, quantitatively, the effects are small. To generate interest rate responses closer to those observed in the data, we need a amplification mechanism. A natural candidate is to introduce a second role for riskess assets and providers of liquidity. For liquidity to be valuable, we need to introduce liquidity constraints.

We embed the liquidity constraint into our baseline framework. Since our innovation - and primary contribution - is to introduce real-time estimation of beliefs as a source of endogenous persistence, we model liquidity constraints in a stylized but tractable specification: Firms have access to a productive opportunity, but require liquidity in the form of pledgeable collateral in order to exploit it. As in Ricardo Lagos (2005), capital and riskless government bonds both can be pledged, albeit to different degrees. An increase in tail risk now has an additional effect – it reduces the liquidity value of capital, increasing the demand for an alternate source of liquidity, namely riskless government bonds. This amplifies the interest rate response of belief revisions.

Preferences and technology: The economy is populated by a unit measure of firms, each endowed with the same Cobb-Douglas production function as the representative firm in the model in Section 2. They are subject to the same aggregate capital quality shocks described in that section.

At the beginning of the period, a fraction ψ of firms have access to a project – which pays $G(X_t)$ for an investment of X_t (both denominated in the single consumption/investment good). The function G is assumed to be strictly increasing and concave, which implies that the net surplus from the project, namely $G(X) - X$ has a unique maximum at X^*. In the absence of other constraints, therefore, every firm presented with this opportunity will invest X^*. However, the firm faces a liquidity constraint

$$X_t \leq B + D + \eta K_t$$

In other words, the investment in the project cannot exceed the sum of pledgeable collateral – which comprises a fraction η of its effective capital K_t and the value of its liquid assets (i.e.

government bonds) B_t – and a debt limit, denoted D. Therefore,

$$X_t = \min\left(X^*, B_t + D + \eta K_t\right)$$

After this stage, both firms with access to the project and those produce consumption goods using capital and labor as before.

The maximization problem of each firm is

$$\begin{aligned} V\left(K_t, B_t\right) &= \max_{X_t, N_t, \hat{k}_{t+1}} \quad \psi\left[G\left(X_t\right) - X_t\right] + F\left(K_t, N_t\right) - W_t N_t + K_t\left(1-\delta\right) + B_t \\ &\qquad -P_t B_{t+1} - \hat{K}_{t+1} + \beta \mathbb{E} M_t V\left(K_{t+1}, B_{t+1}\right) \end{aligned}$$

$$\begin{aligned} s.t. \quad X_t &\leq B_t + \eta K_t + D \\ K_{t+1} &= \phi_{t+1} \hat{K}_{t+1} \end{aligned}$$

Each firm faces the same problem because all firms are identical ex-ante.

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

1. Firms enter the period with a capital stock $\hat{K}_t$ and liquid assets B_t.
2. The aggregate capital quality shock ϕ_t and the investment opportunity shocks are realized. Those with the opportunity choose X_t subject to the liquidity constraint.
3. Firms choose labor and production takes place. Firms choose capital and liquid asset positions for $t+1$.

5.1 Characterization and solution

The equilibrium of the economic model is a solution to a set of non-linear equations. In this section, we derive the equations characterize the equilibrium 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.

The key optimality conditions of the firm's problem are

$$\begin{aligned} 1 &= \beta \mathbb{E}_t\left\{M_{t+1}\phi_{t+1}\left[F_1\left(K_{t+1}, N_{t+1}\right) + 1 - \delta + \eta\mu_t\right]\right\} \\ P_t &= \beta \mathbb{E}_t\left\{M_{t+1}\left[1+\mu_t\right]\right\} \\ \mu_t &= \psi(G'\left(X_t\right) - 1) \end{aligned}$$

where μ_t is the Lagrange multiplier on the liquidity constraint. The first two equations are the Euler equation for capital and liquid assets respectively. The liquidity value of these assets is reflected in the last term on the right hand side (the one with the μ_t).

5.2 Calibration

This section calibrates the model to examine the quantitative predictions. The utility and production functions are

$$\begin{aligned} U(C,N) &= \frac{\left(C - \pi \frac{N^{1+\gamma}}{1+\gamma}\right)^{1-\sigma}}{1-\sigma}, \\ F(K,N) &= K^{\alpha} N^{1-\alpha}, \\ G(X) &= 2\zeta\sqrt{X}. \end{aligned}$$

As before, a period is interpreted as a year. We choose the discount factor $\beta = 0.95$ and depreciation $\delta = 0.06$ (consistent with the reported depreciation rates in the Flow of Funds data) to target a steady state capital-output ratio of 3.5 (this is taken from Cooley and Prescott (1995)). The share of capital in the production, is 0.40, which is also taken from Cooley and Prescott (1995). The labor supply parameter, γ, is set to 0.5, in line with Midrigan and Philippon (2011), corresponding to a Frisch elasticity of 2. The labor disutility parameter π is normalized to 1. The collateral parameter η, is set to match a measured return on capital $F_1(K,N) + 1 - \delta = 1.03$. The sum of liquid assets and debt $B + D$ are set 75% of capital. The parameter ζ is set equal to 4.9 to generate a risk-free rate of 2% (in the stochastic steady state associated with $\hat{g}_{2007}$), and ψ is set equal to 0.25. Finally, the risk aversion σ is set equal to 5 [10]

5.3 Results

Our main goal in this section is to quantify the size and persistence of the macroeconomic response to a large but transitory shock ϕ_t and explore how belief updating alters the interest rate response in the long run.

To assess the economic effect of belief revisions following the financial crisis, we begin by estimating $\hat{g}_{2007}$ using historical data on ϕ_t from 1950-2007, measured using the strategy outlined in Section 3 and computing the stochastic steady state associated with this distribution.[11] Then,

[10] In Section 5.4 we show some robustness with respect to risk aversion and the collateral parameter.

[11] 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.

Parameter	Value	Description
Preferences:		
β	0.95	Discount factor
γ	0.50	1/Frisch elasticity
π	1	Labor disutility
σ	5	Risk aversion
ζ	4.9	Risk aversion
Technology:		
α	0.40	Capital share
δ	0.06	Depreciation rate
Liquidity:		
$B + D$	7.0	Supply of liquid assets
η	0.70	Collateral parameter

Table 3: **Parameters**

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 by drawing thousands of sequences from the revised estimate $\hat{g}_{2009}$. For each of these sequences of shocks, we compute the sequence of equilibrium prices and quantities. Finally, we average over all these random outcomes and record the mean response. Figure 4 plots the mean future path of the real interest rate. The model generates prolonged low interest rates, where the economy (on average) never recovers from the negative shocks in 2008-'09. Instead, all aggregate variables move towards the new, lower (stochastic) steady state.

It is important to note that the results in Figure 4 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, none of which contain a crisis, the shock will still be persistent. But in the long-run, outcomes may resemble 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.

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

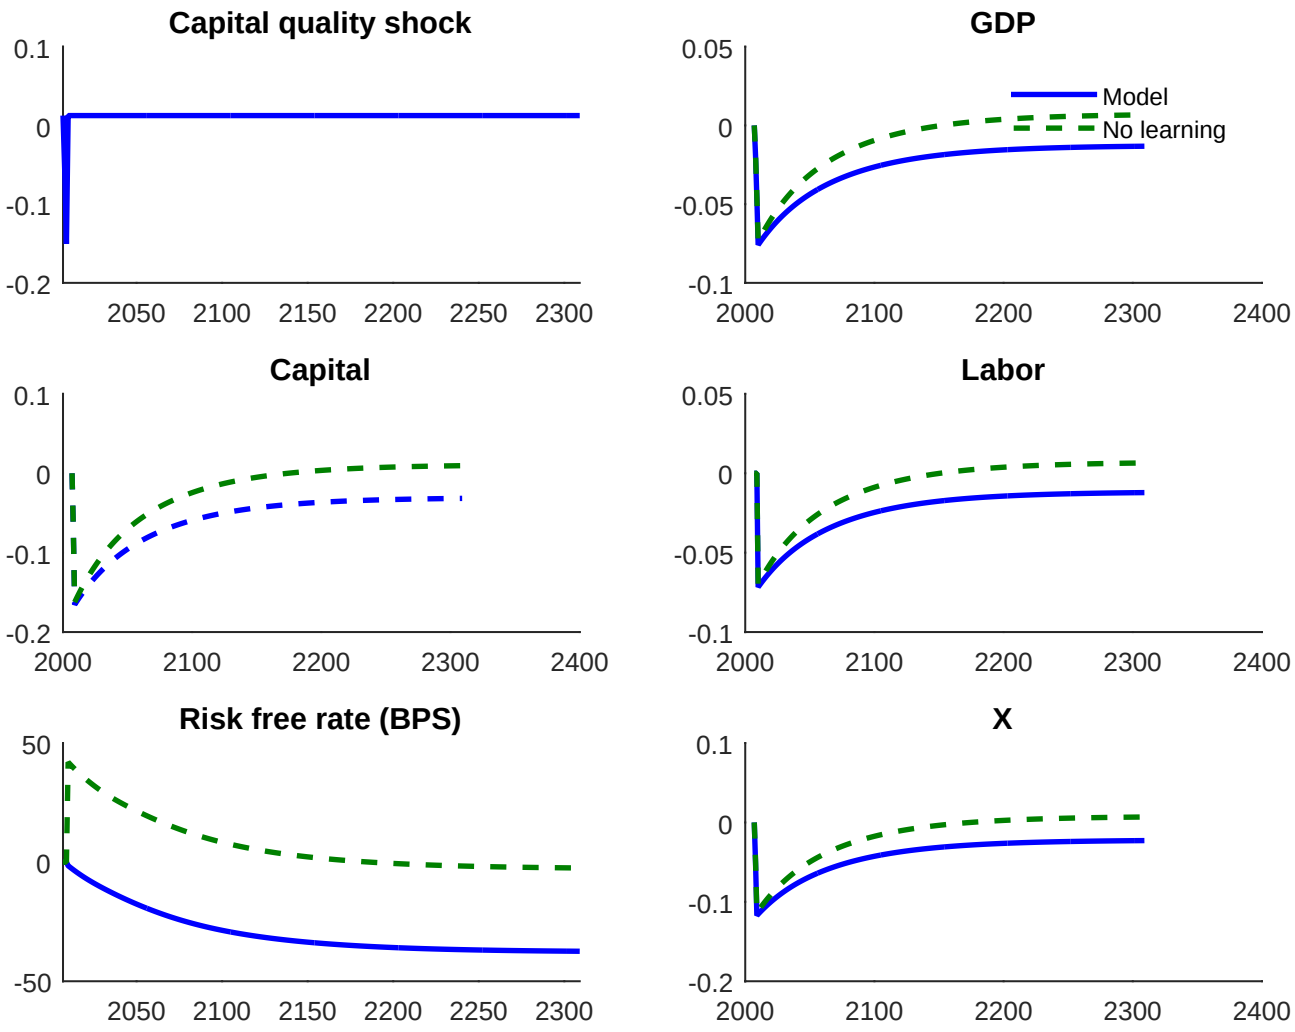


Figure 4: **Large negative shocks create extremely persistent interest rate declines.** *Solid line shows the change in aggregates (relative to the stochastic steady state associated with $\hat{g}_{2007}$). Dashed line (no learning) is an identical model where agents believe that shocks are drawn from $\hat{g}_{2009}$ and never revise those beliefs.*

$\hat{g}_{2009}$ with certainty, from the very beginning (dashed line in Figure 4). 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 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 causes the real rate to surge and then recover. The rise in the interest rate arises because, as the economy is recovering back to the

previous steady state, there is a lower demand for debt.[12] This shows that learning is what generates long-lived reductions in economic activity.

5.4 Robustness

Table 4 shows the change in the risk free rate and output under different values of risk aversion (σ) and collateral (η). When η increases from 0 to 0.7, capital is not only valuable for production but also as collateral. Hence, after a change in beliefs, capital (and therefore output) is reduced by a lower amount when $\eta = 0.7$. However, when capital serves as collateral, the effects on the risk free rate are larger. Intuitively, as there is less capital in the economy, the demand for risk free assets increases, and the risk free rate diminish.

	Long-run effects	
	$\eta = 0$	$\eta = 0.7$
Risk free rate		
$\sigma = 1$	-4 bps	-25 bps
$\sigma = 2$	-8 bps	-28 bps
$\sigma = 5$	-25 bps	-40 bps
Output		
$\sigma = 1$	-6.37%	-3.61%
$\sigma = 2$	-5.96%	-3.29%
$\sigma = 5$	-3.91%	-2.00%

Table 4: **Robustness: Different risk aversion and collateral parameters.**
The table shows the long-run effects, i.e., the change in the risk free rate and output in the stochastic steady state associated with $\hat{g}_{2009}$) relative to $\hat{g}_{2007}$.

6 Conclusion

No one knows the true distribution of shocks to the economy. Economists typically assume that agents in their models do know this distribution as a way to discipline beliefs. But assuming that agents do the same kind of real-time estimation that an econometrician would do is equally disciplined and more plausible. For many applications, assuming full knowledge has little effect on outcomes and offers tractability. But for outcomes that are sensitive to tail probabilities, the difference between knowing these probabilities and estimating them with real-time data can be large. The estimation error can be volatile and can introduce new, persistent dynamics into a

[12]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).

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 face liquidity constraints, the value of riskless, liquid assets is sensitive to tail risk. Tail events are states when output is low and liquid assets are extremely valuable. Therefore, the yield on a liquid asset depends precisely on the probability of a tail event, which is hard to estimate and whose estimates fluctuate greatly. When a shock arrives that is a negative outlier, tail risk surges, and the value of liquid assets rises. The rise in the value of the asset shows up as a fall in its yield.

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 persistently low interest rates took hold because, after seeing how fragile our financial sector is, market participants will never think about tail risk in the same way again.

References

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

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

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

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).

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

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

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

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

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

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.

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

Kozlowski, J., L. Veldkamp, and V. Venkateswaran (2015). The tail that wags the economy: Belief-driven business cycles and persistent stagnation. New York University working paper.

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.

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.

Ricardo Lagos, R. W. (2005). A unified framework for monetary theory and policy analysis. *Journal of Political Economy 113*(3), 463–484.

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

Appendix

A Optimality conditions from firm's problem

Envelope conditions are

$$\begin{aligned} V_1(K_t, B_t) &= F_1(K_t, N_t) + 1 - \delta + \mu_t \eta \\ &= F_1(K_t, N_t) + 1 - \delta + \psi\eta(G'(X_t) - 1) \\ V_2(K_t, B_t) &= 1 + \psi(G'(X_t) - 1) \end{aligned}$$

Aggregate consumption and labor supply are

$$\begin{aligned} C_t &= \psi(G(X_t) - X_t) + F(K_t, N_t) + K_t(1-\delta) - \hat{K}_{t+1} \\ M_{t+1} &= \frac{U_1(C_{t+1}, N_{t+1})}{U_1(C_t, N_t)} \\ U_1(C_t, N_t) &= -F_2(K_t, N_t) U_2(C_t, N_t) \end{aligned}$$

Functional forms

$$G(X) = 2\zeta\sqrt{X} \qquad \Rightarrow \qquad \frac{\zeta}{\sqrt{X^*}} = 1$$

Steady state:

$$\begin{aligned} P &= \beta\left(1 + \psi\left(\frac{\zeta}{\sqrt{\bar{X}}} - 1\right)\right) \\ 1 &= \beta\left[F_1(\bar{K}) + 1 - \delta + \psi\eta\left(\frac{\zeta}{\sqrt{\bar{X}}} - 1\right)\right] \\ \bar{X} &= \bar{B} + \eta\bar{K} + \bar{D} \\ C &= 2\zeta\sqrt{\bar{X}} - \bar{X} + F(\bar{K}) - \delta\bar{K} \end{aligned}$$

Calibration: $(\psi, \zeta, \eta, \bar{B} + \bar{D})$. Targets

$$P = 0.98 \qquad lev = \frac{\bar{B} + \bar{D}}{\bar{K}} = 0.75 \qquad R_k = F_1(\bar{K}) + 1 - \delta = 1.05 \qquad \psi = 0.10$$

Then,

1. EE for Bonds

$$\begin{aligned} P &= \beta\left(1 + \psi\left(\frac{\zeta}{\sqrt{\bar{X}}} - 1\right)\right) \\ \frac{P}{\beta} - 1 &= \psi\left(\frac{\zeta}{\sqrt{\bar{X}}} - 1\right) \end{aligned}$$

2. Use previous EQ + EE K to solve for η

$$\begin{aligned} 1 &= \beta\left[R_k + \eta\psi\left(\frac{\zeta}{\sqrt{\bar{X}}} - 1\right)\right] = \beta\left[R_k + \eta\left(\frac{P}{\beta} - 1\right)\right] \\ \eta &= \frac{\frac{1}{\beta} - R_k}{\left(\frac{P}{\beta} - 1\right)} = \frac{1 - \beta R_k}{P - \beta} \end{aligned}$$

3. Use EE K to solve for $\bar{K}$

$$\begin{aligned} 1 &= \beta\left[F_1\left(\bar{K}\right) + 1 - \delta + \eta\psi\left(\frac{\zeta}{\sqrt{\bar{X}}} - 1\right)\right] \\ 1 &= \beta\left[F_1\left(\bar{K}\right) + 1 - \delta + \eta\left(\frac{P}{\beta} - 1\right)\right] \\ \alpha\bar{K}^{\alpha-1} &= \frac{1}{\beta} - 1 + \delta - \eta\left(\frac{P}{\beta} - 1\right) \\ F_1\left(\bar{K}\right) &= \left(\frac{1}{\alpha}\left(\frac{1}{\beta} - 1 + \delta - \eta\left(\frac{P}{\beta} - 1\right)\right)\right)^{\frac{1}{\alpha-1}} \end{aligned}$$

4. Use constraint to solve for $\bar{X}$

$$\bar{X} = \left(\frac{\bar{B} + \bar{D}}{\bar{K}} + \eta\right)\bar{K}$$

5. Use leverage to solve for $\bar{B} + \bar{D}$

$$\bar{B} + \bar{D} = lev\bar{K}$$

6. Use EE Bonds to solve for ζ

$$\zeta = \left(1 + \frac{1}{\psi}\left(\frac{P}{\beta} - 1\right)\right)\sqrt{\bar{X}}$$