Federal Reserve Bank of Chicago

Risk Premia at the ZLB: A Macroeconomic Interpretation

François Gourio and Phuong Ngo

January 2, 2020

WP 2020-01

https://doi.org/10.21033/wp-2020-01

Risk Premia at the ZLB: A Macroeconomic Interpretation*

François Gourio† and Phuong Ngo‡

January 2, 2020

Abstract

Historically, inflation is negatively correlated with stock returns, leading investors to fear inflation. We document using a variety of measures that this association became positive in the U.S. during the 2008-2015 period. We then show how an off-the-shelf New Keynesian model can reproduce this change of association due to the binding zero lower bound (ZLB) on short-term nominal interest rates during this period: in the model, demand shocks become more important when the ZLB binds because the central bank cannot respond as effectively as when interest rates are positive. This changing correlation in turn reduces the term premium, and hence contributes to explaining the decline in long-term interest rates. We use the model to evaluate this mechanism quantitatively. Our results shed light on the validity of the New Keynesian ZLB model, a cornerstone of modern macroeconomic theory.

JEL classification: C61, E31, E52, E62.
Keywords: zero lower bound, liquidity trap, stock market, inflation premia, term premia, risk premia.

*We are grateful to Fernando Alvarez, Stefania D'Amico, Gadi Barlevy, Robert Barsky, Marco Bassetto, Jeff Campbell, Jesus Fernandez-Villaverde, Jon Steinsson, Andrea Tambalotti, Pietro Veronesi, many other colleagues, and seminar or conference participants at the Federal Reserve Banks of Boston, Chicago, Cleveland, and the Board of Governors, Duke University, University of Southern California, University of Montreal, the University of British Columbia, Chicago Booth, the Ohio State University, Ohio University, Vanderbilt University, London School of Economics, International Monetary Fund, the 2016 NBER Summer Institute, the 2016 NEOEW, 2016 SED Annual Meeting, the 2017 WFA and MFA Annual Meetings, and especially to our discussants Jordi Gali, Carolin Pflueger, Dongho Song and Oreste Tristani for helpful comments and suggestions. In addition, Phuong Ngo gratefully acknowledges the supports from the Office of Research at Cleveland State University and the Ohio Supercomputer Center (1987). The views expressed here are those of the authors and do not necessarily represent those of the Federal Reserve Bank of Chicago or the Federal Reserve System.

†Federal Reserve Bank of Chicago; Email: francois.gourio@chi.frb.org.
‡Cleveland State University; Email p.ngo@csuohio.edu.

1 Introduction

The relation between inflation and economic activity is controversial, as illustrated by the debate about the empirical relevance of the Phillips curve. The purpose of this paper is to use financial markets data to shed light on this relation, and in particular to document and understand the changes in the association of inflation and stock returns, to use this as a test of the standard New Keynesian model, and to study the implications for asset pricing.

Our starting point is the observation that there has been a significant change in the association of inflation and the market return after 2008. Historically, high inflation is associated with low stock returns, as documented in a long literature dating back at least to Fama and Schwert (1977) and Modigliani and Cohn (1979). We provide a variety of evidence that this association changed during the period 2008-2015, during which inflation and stock returns are positively, rather than negatively, associated. As a simple illustration, figure 1 depicts the strong positive correlation between the S&P 500 index and the 10-year breakeven (the difference between the yield of a 10-year Treasury bond and a 10-year Treasury inflation-indexed bonds, which is often used as a proxy for expected inflation) between 2009 and 2013. During this period, movements in stock prices - which often reflected news about the economic recovery - were associated with movements in inflation breakevens. Hence, it appears that financial markets viewed signals about the economy and about future inflation as positively correlated - a significant deviation from the historical record.

Why is the association of stock returns and inflation of interest? First, from an asset pricing perspective, this association measures whether investors should be averse to inflation. When inflation is countercyclical,[1] investors require a premium to hold nominal assets. This premium is then embedded in nominal interest rates. If inflation becomes less countercyclical, or even procyclical, this premium is reduced or can even become negative, reducing the level of long-term interest rates, which might in turn affect the economy. Second, from a macroeconomic perspective, it is tempting to read the association of stock returns and inflation depicted above as a vindication of sorts for the Phillips curve. Even though the association of inflation and economic activity is weak in the data, financial markets reveal that expectations of future inflation (as proxied by the breakeven) are strongly correlated with expectations of future output (as proxied by the stock market). This would seem to be an informative data point for macroeconomists.

[1] By countercyclical, we mean here that inflation is high in "bad times", i.e. states with high discount factor, e.g. a low market return in a CAPM model, or a high marginal utility of consumption (low consumption) in a consumption-based model.

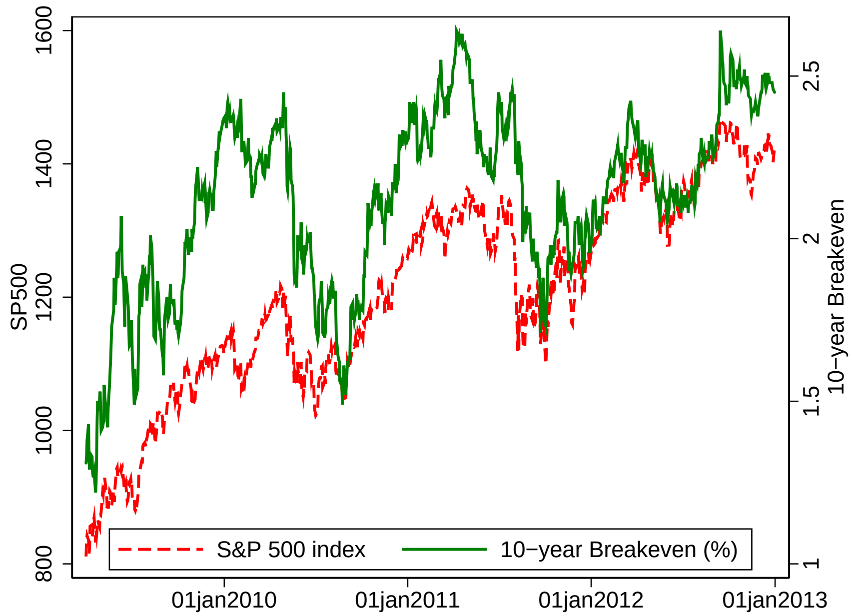

Figure 1: S&P 500 index level (left scale) and 10-year inflation breakeven (right scale), from July 2009 to December 2012.

Our aim in this paper is to explain the changing association and to draw out its implications for asset pricing and for macroeconomics and in particular for the standard New Keynesian model, a cornerstone of modern macroeconomic theory. We first show how this off-the-shelf model can explain the changing correlation owing to the zero lower bound (ZLB) constraint on monetary policy. The ZLB constraint creates a significant nonlinearity which can generate precisely such a change in correlation. This result stands in stark contrast with standard macroeconomic models that are typically log-linear, and hence have a constant covariance structure. The key economic mechanism driving this change in correlation is that the response to a macroeconomic shock depends on whether the ZLB binds. In our model, the economy is subject to both supply (productivity) and demand (liquidity preference) shocks. In normal times (e.g. in steady-state, far from the ZLB), positive supply shocks lead to high output and stock returns, and low inflation, while positive demand shocks lead to high output, stock returns, and inflation. This pattern is consistent with simple aggregate supply - aggregate demand framework (AS/AD) and also emerges in almost all New Keynesian models. Quantitatively, in normal times, supply shocks dominate, generating an overall negative covariance of inflation and stock returns, and hence a positive inflation risk premium, and a positive term premium, in line with the data pre-2008.

However, the propagation of shocks becomes notably different once the ZLB binds. Focus for simplicity here on demand shocks. Such a shock has relatively little effect on inflation or economic activity if the ZLB does not bind, because the central bank can offset demand fluctuations by adjusting the short-term nominal interest rate. But the same demand shocks may have large effects if the ZLB binds and the central bank cannot respond. As a result, the importance of demand shocks for stock returns and inflation increases, increasing their covariance and shifting it from negative to less negative, or even positive. This in turn implies a lowering of inflation risk premia as the economy becomes closer to the ZLB, which is consistent with evidence from affine term structure models.[2]

This lowering of term premia has a number of implications. First, lower term premia imply lower long-term interest rates, hence the model helps explain why long-term interest rates have remained so low since 2008. However, it turns out that in our model, this effect is quantitatively limited, explaining only about 10% of the change in interest rates. Second, economists and policymakers often use inflation breakevens as a measure of expected inflation. It is well understood that inflation compensation may differ from expected inflation due to risk or liquidity premia; but the magnitude and even the sign of this adjustment are controversial. Our model argues that breakevens *underestimate* expected inflation when the economy operates close to the ZLB, but *overestimate* expected inflation when the economy is far from the ZLB. Third, our model suggests that an upturn in the economy or in inflation may lead to a significant increase in interest rates because these risk premia might increase as the ZLB becomes less of a constraint.

Finally, our analysis is a test of the widely-used ZLB New Keynesian macroeconomic model. In our simulations, we find that while the model does well qualitatively, its quantitative performance is more mixed. On the positive side, the model easily replicates the magnitude of the changing association of inflation and stock returns. On the negative side, term premia remain too small due to the relatively transitory nature of shocks in the model. Moreover, the model implies too large a deflation during the ZLB period - a well-known limitation of this framework. As a result, the implications for real interest rates are also at odds with the data. Overall, our results outline some important challenges for the New Keynesian ZLB model, which result from the novel evidence that we use to confront the model.

To sum up, the paper makes three main contributions. First, we document the change in the association between inflation and stock returns during the 2008-2015 period where the ZLB

[2] For some prominent models, see Kim and Wright (2005), Adrian et al. (2013), DAmico et al. (2018), Ajello et al. (2014).

was binding in the United States. Second, we show how a standard macroeconomic model can be at least qualitatively consistent both with the negative association pre-2008 and the positive association post-2008, as explained above, i.e. the model generates endogenously a time-varying covariance that matches the data. Third, we study the implications of this model for asset prices and for macroeconomics, and discuss the successes but also the limitations of this framework.

The paper is organized as follows. The rest of the introduction reviews briefly the related literature. Section 2 studies a simple example that illustrates what determines the covariance of inflation and stock returns and why this covariance matters for asset prices. Section 3 presents empirical evidence that the link between the inflation and stock returns changed after 2008. Section 4 introduces a stylized DSGE model and Section 5 studies it quantitatively. Section 6 concludes. An online appendix describes additional empirical results and our numerical method.

Related Literature

Our paper is related to several strands of literature. First, there is a large macro-finance literature that studies the term structure of interest rates and in particular the inflation risk premium, including affine models and their extensions such as Ang et al. (2008), Hordahl and Tristani (2014), or Roussellet (2018). The estimates of risk premia from these models are corroborated by survey measures in Wright (2011) or Breach et al. (2016). Most term structure models also find that the inflation risk premium has been low post 2008. Fleckenstein et al. (2014) and Fleckenstein et al. (2017) also study the pricing of TIPS and deflation risk. Representative agent endowment economy models have also been proposed, notably Piazzesi and Schneider (2006) and Bansal and Shaliastovich (2013). The underlying logic of how risk premia are determined is similar to our paper (and is discussed in section 2), but our contribution relative to these papers is to study the *sources* of the correlations between inflation and growth that are taken as primitives in these studies. David and Veronesi (2013) also study the changes in regimes with different correlations of inflation and asset prices.

Second, several authors have proposed DSGE production models with nominal rigidities that attempt to replicate various features of asset prices. Key contributions include Rudebusch and Swanson (2008), Rudebusch and Swanson (2008), Li and Palomino (2014), Christiano et al. (2010), Palomino (2012), and Swanson (2015a). Especially close to our work are the recent papers by Campbell et al. (2014), Campbell et al. (2009), Song (2017), and Branger et al. (2016)

which emphasize structural breaks in monetary policy rules and how these affect asset prices and their correlations since the late 1990s and early 2000s. Our contribution relative to all these papers is to introduce the ZLB and to focus on the recent changes since the Great Recession started. The contemporaneous studies by Nakata and Tanaka (2016), Datta et al. (2018), and Bilal (2017) are also closely related, with a somewhat similar message but differences in empirical work and theoretical model.

Third, our paper relates to the vast macroeconomic literature on the effects of the zero lower bound. Seminal contributions include Krugman (1998) and Eggertsson and Woodford (2003) as well as Eggertsson and Krugman (2012) and Nakata (2017).[3] While the relevance of the ZLB is largely acknowledged by macroeconomists and policymakers, the importance of New Keynesian model mechanisms remains in dispute (see Wieland (2014) for a powerful critique). Our paper adds novel data to test the New Keynesian model mechanism. Debortoli et al. (2019) in follow-up work study structural breaks in macroeconomic dynamics post ZLB, and finds little evidence of a change. In contrast, our paper uses financial markets data which may be more informative, as we discuss below in more detail. Another critique of the ZLB analysis is the lack of response of inflation during the Great Recession (e.g. Del Negro et al. (2015), Bianchi and Melosi (2017), Gilchrist et al. (2017)), which will also feature in our results.

Finally, the broader question of the relation between stock prices and inflation has long a long history dating back at least to Fama and Schwert (1977) who showed that stocks appeared to be affected negatively by inflation, a result widely viewed as "puzzling" since stocks are claims to real assets. Modigliani and Cohn (1979) argued that investors suffered from money illusion. Boudoukh and Richardson (1993) and Campbell and Vuolteenaho (2004) revisited the empirical correlation. Marshall (1992) is an early structural model of the relationship between inflation and stock returns. Duarte (2013) also emphasizes the change in correlation over time and studies how inflation affects the cross-section of stock returns.[4] Our empirical analysis uses high-frequency data on the response of stock prices to macroeconomic data releases, as in Boyd et al. (2005), Faust et al. (2007), Law et al. (2018), and Ai and Bansal (2018).

[3] In particular, our nonlinear method is related to the contributions of Fernandez-Villaverde et al. (2015), Ngo (2014), Gust et al. (2017), and Miao and Ngo (2014).

[4] Also related is the research on the effect of monetary policy on stock prices. Bernanke and Kuttner (2005) and Rigobon and Sack (2004) showed that monetary policy shocks have a large positive effect on stock prices. Gorodnichenko and Weber (2016) and Weber (2015) study cross-sectional differences in the heterogeneity in price flexibility and demonstrate that it affects the responses of individual stocks to monetary policy shocks.

2 Sources and implications of the inflation-stock return covariance

This section uses a simple representative agent, endowment economy model to explain why the inflation-stock return covariance might change, and why it affects asset prices. Suppose that the representative consumer has expected utility with constant relative risk aversion:

$$E\sum_{t=0}^{\infty}\beta^t\frac{C_t^{1-\gamma}}{1-\gamma},$$

and that consumption growth and inflation are jointly log-normally distributed and *iid*, that is

$$\begin{pmatrix}\Delta c_{t+1}\\ \pi_{t+1}\end{pmatrix}\sim N\left(\begin{pmatrix}\mu_c\\ \mu_p\end{pmatrix},\begin{pmatrix}\sigma_c^2 & \rho_{c,p}\\ \rho_{c,p} & \sigma_p^2\end{pmatrix}\right),$$

where lowercase letters denote logs, so that $\Delta c_{t+1}=\Delta\log C_{t+1}$ is (log) consumption growth, and log inflation is $\pi_{t+1}=\Delta\log P_{t+1}$, where P_t is the consumer price index. Here the parameters μ_c,μ_p and $\sigma_c,\sigma_p,\rho_{cc,p}$ describe the equilibrium stochastic process followed by consumption growth and inflation. (In section 5, we will endogenize the process for consumption and inflation using a dynamic stochastic general equilibrium model, but for now it is easiest to take this as given.)

The stochastic discount factor used to price real assets is, in log:

$$m_{t+1}=\log M_{t+1}=\log\beta-\gamma\Delta c_{t+1},$$

and the stochastic discount factor used to price nominal assets is:

$$m_{t+1}^{\$}=\log M_{t+1}^{\$}=m_{t+1}-\pi_{t+1}.$$

We model the stock as a claim to levered consumption, that is the log dividend is proportional to log consumption,

$$d_t=\xi c_t,$$

where ξ captures financial or operating leverage, as in Abel (1999). It is easy to solve for the stock price and obtain that the log gross stock return is, up to an unimportant constant,

$$\log R_t^s=\xi\Delta c_t.$$

Hence, the covariance of stock return and inflation is simply proportional to $\rho_{c,p}$:

$$Cov_t\left(\log R_{t+1}^s,\pi_{t+1}\right)=\xi\rho_{c,p}.$$

We first explain why the covariance of stock returns and inflation matters, before explaining why it might change.

2.1 Implications

The covariance of stock returns and inflation is important because it reflects the parameter $\rho_{c,p}$, which is a determinant of the inflation risk premium. The covariance $\rho_{c,p}$, which may be positive or negative, measures the association of inflation and marginal utility (which is consumption growth in this simple example). Intuitively, a positive $\rho_{c,p}$ corresponds to the case where "demand shocks" dominate: low consumption (high marginal utility) is associated with low inflation, while a negative $\rho_{c,p}$ corresponds to the case where "supply shocks" dominate: low consumption (high marginal utility) is associated with high inflation.

To see the connection to the inflation risk premium, we now calculate the yield on nominal and real bonds. For simplicity, we will focus on one-period bonds; one may think of the time period as being long, for instance a few years. The log of the gross real risk-free rate is

$$\begin{aligned}\log R^f_{t+1} &= -\log E_t\left(M_{t+1}\right), \\ &= -\log\beta + \gamma\mu_c - \frac{\gamma^2\sigma_c^2}{2}.\end{aligned}$$

This familiar formula decomposes the riskless rate into impatience, intertemporal substitution, and precautionary savings. The log nominal risk-free rate is, on the other hand,

$$\begin{aligned}\log R^{f,\$}_{t+1} &= -\log E_t\left(M^{\$}_{t+1}\right), \\ &= \log R^f_{t+1} + \mu_p - \frac{1}{2}\sigma_p^2 - \gamma\rho_{c,p}\end{aligned}$$

The nominal rate equals the real rate plus an inflation compensation term that reflects expected inflation and inflation risk. It is useful to separate this inflation compensation and define the breakeven rate, i.e. the difference in the (log) yields of these two bonds:

$$\begin{aligned}\log BE_t &\equiv \log R^{\$}_{t+1} - \log R_{t+1} \\ &= \mu_p - \frac{1}{2}\sigma_p^2 - \gamma\rho_{c,p}.\end{aligned}$$

This shows that the breakeven rate is the sum of expected (log) inflation, a Jensen adjustment,[5] and a risk premium term which equals risk aversion γ multiplied by the covariance $\rho_{c,p}$. (This term will be large provided risk aversion is large, which is necessary to replicate asset prices.)

Intuitively, if the covariance $\rho_{c,p} < 0$, supply shocks dominate, and breakevens overestimate expected inflation. Nominal bonds are risky assets, since their real payoff is low in states of

[5] The source of this term is that the real payoff of a nominal bond depends inversely on inflation. Consequently, higher uncertainty about inflation leads to higher expected payoffs. This term is typically small.

the world where inflation is high, which on average coincide with low consumption growth and high marginal utility. Hence, agents require a premium to hold nominal bonds, so the nominal yield is higher than it would be under risk-neutrality. On the other hand, if $\rho_{c,p} > 0$, demand shocks dominate, inflation is a hedge, breakevens underestimate inflation, and nominal yields are lower.

These formulas explains why the covariance term $\rho_{c,p}$ plays a crucial role, and why it can be calculated using the covariance of stock returns and inflation. We will argue that $\rho_{c,p}$ was negative in the pre-2008 "regime" and became less negative, and perhaps even positive, in the 2008-2015 ZLB "regime".

In this simple model, this term can be calculated using the covariance of consumption growth and inflation as well as stock returns and inflation. We believe there are several good reasons to focus on the later. First, from a measurement perspective, consumption is noisy, and measured only at a relatively slow frequency, whereas stock returns can be measured quickly without error. This is especially useful when considering short time series and looking for a regime shift. Second, in most asset pricing models, stock returns reflect shocks that affect marginal utility, even if aggregate nondurable consumption is not the main factor affecting marginal utility (i.e. the consumption CAPM is wrong). But we will also discuss the measurement of the covariance of economic activity and inflation later.

2.2 Sources

We have discussed why the covariance of stock returns and inflation matters; now we discuss what determines this covariance. To do so requires going one step further and decomposing stock returns and inflation into their responses to fundamental macroeconomic shocks. For simplicity, suppose there are two fundamental shocks, "demand" and "supply", and that inflation goes up with the "demand" shock, but down with the "supply" shock. We can write this as:

$$R^s_{t+1} = \lambda_{r,d}\varepsilon_{d,t+1} + \lambda_{r,s}\varepsilon_{s,t+1},$$
$$\pi_{t+1} = \lambda_{\pi,d}\varepsilon_{d,t+1} + \lambda_{\pi,s}\varepsilon_{s,t+1},$$

where $\varepsilon_{d,t+1}, \varepsilon_{s,t+1}$ are independent shocks with variances σ_d^2, σ_s^2 respectively, and the parameters $\lambda_{c,d}, \lambda_{c,s}, \lambda_{\pi,d}, \lambda_{\pi,s}$ reflect the responses to these shocks (which depend on structural parameters in the DSGE model, but here are written as reduced form parameters). We have $\lambda_{r,d} > 0$,and $\lambda_{r,s} > 0$ (a normalization) and $\lambda_{\pi,d} > 0$, $\lambda_{\pi,s} < 0$ from the discussion above. The covariance can be easily calculated as

$$Cov_t(R^s_{t+1}, \pi_{t+1}) = \underbrace{\lambda_{r,d}\lambda_{\pi,d}}_{>0}\sigma_d^2 + \underbrace{\lambda_{r,s}\lambda_{\pi,s}}_{<0}\sigma_s^2.$$

The first term, which is positive, reflects that demand shocks move stock returns and inflation in the same direction, while the second term reflects that supply shocks move them in opposite direction. The overall covariance, then, depends on the variance of the shocks, as well as the coefficients λ's which reflect precisely how much do stock returns and inflation respond to a shock. In the next paper, we will argue that the zero lower bound leads to an increase in $\lambda_{r,d}$ and $\lambda_{\pi,d}$, leading the covariance to increase.[6]

3 Changes in the relation between stock prices and inflation

This section documents that the association between inflation and stock returns became more positive around 2008. We use four different approaches to measure the association: first, we use high-frequency financial markets data on inflation compensation, i.e. measures of inflation expectations embedded in asset prices; second, we review the response of stock returns to inflation data releases; third, we construct from the cross-section of U.S. stocks a portfolio that "mimics" inflation and study its correlation with the aggregate stock market; and fourth, we look at the simple correlation between realized monthly inflation and stock returns. We also present briefly some evidence on the association of economic activity and inflation.

When exactly does the association of stock returns and inflation change? This is a complicated question as there are secular changes in this relation (as argued for instance in Campbell et al. (2014)), as well as business-cycle frequency changes, such as the one we focus on. For transparency and simplicity, all our measures use the same cutoff date, namely July 2008. This cutoff date is the one that is most often selected by formal statistical break tests, as we show in a last subsection. It is also suggested by intuitive rolling regressions, as we show in the online appendix. Moreover, it approximately corresponds to the start of the most serious phase of the financial crisis, during which the Federal Reserve quickly lowered the interest rate to hit the "zero lower bound". We end the sample in December 2015, when the Federal Funds rate eventually "lifts off" from zero.[7]

[6] Another possibility would be that agents revised their estimates of the variances of the fundamental shocks.

[7] One could argue for an earlier end of sample if the Fed was effectively keeping interest rates lower than normal for longer (as in Eggertsson and Woodford (2003)). Ending the sample in 2013 or 2014 has no material effect on our results. Changing the start date to some other date in 2008 or even early 2009 has relatively little effect on our results as well, except for the results of section 3.2.

3.1 Inflation compensation

Generally, inflation compensation refers to the extra price required by the buyer of an asset that is exposed to inflation over one that is not. The most common measure (the "breakeven inflation") is the difference between the yield on a nominal Treasury security and the yield on an identical maturity inflation-protected security (TIPS). We start by illustrating graphically how the correlation of inflation compensation and stock prices changes after 2008, before presenting statistical evidence.

Figure 2 depicts the daily return on the S&P 500 index against the daily changes in the 10-year breakeven. (For easier visualization, we present our data in ventiles (20 groups) using a standard bin-scatter; the fitted lines, as well as the regressions reported below use the underlying daily data.) The left panel (A) shows a weak correlation in the 2003-2008 sample,[8] while in the right panel (B) the correlation becomes quite strong after 2008.

One might think that these daily swings reflect changing demand for nominal Treasury securities, which have a special liquidity. To assess this, in panels (C) and (D) we replace the breakeven with inflation swaps, which are derivatives with (nearly) identical payoffs as breakevens, but which do not have the specific liquidity attributes of Treasuries. The patterns (estimated starting in January 2004 when our inflation swap data become available) are qualitatively and quantitatively very similar. Panels (E) and (F) show the result when we use an "inflation portfolio" instead as measure of inflation compensation; we discuss the construction of this portfolio in section 3.4 below.

Another possible interpretation is that daily data reflects "market sentiment" rather than real news. There is indeed some evidence that the response of financial markets to macroeconomic news is stronger when cumulated over a few weeks (Altavilla et al. (2017)). Figure 2 in the appendix depicts the correlation between 20-day returns in the S&P 500 index against 20-day changes in 10-year breakeven and inflation swaps. The change in the relation between the two subsamples is even more pronounced using this slightly lower frequency.

To measure the association and establish its statistical significance, we run the following daily regression:

$$\Delta IC_t = \beta_0 + \beta_1 D_{t\geq 2008:7} + \beta_2 R_t^s + \beta_3 D_{t\geq 2008:7} \times R_t^s + \varepsilon_t,$$

where ΔIC_t is the change in the measure of inflation compensation; $D_{t\geq 2008:7}$ is a dummy equal to 1 after July 1st, 2008; R_t^s is the S&P5 00 return; and we estimate the relation over the sample

[8] TIPS were introduced in 1997 but liquidity remained limited until the early 2000s. We follow most of the literature and start our analysis in January 2003.

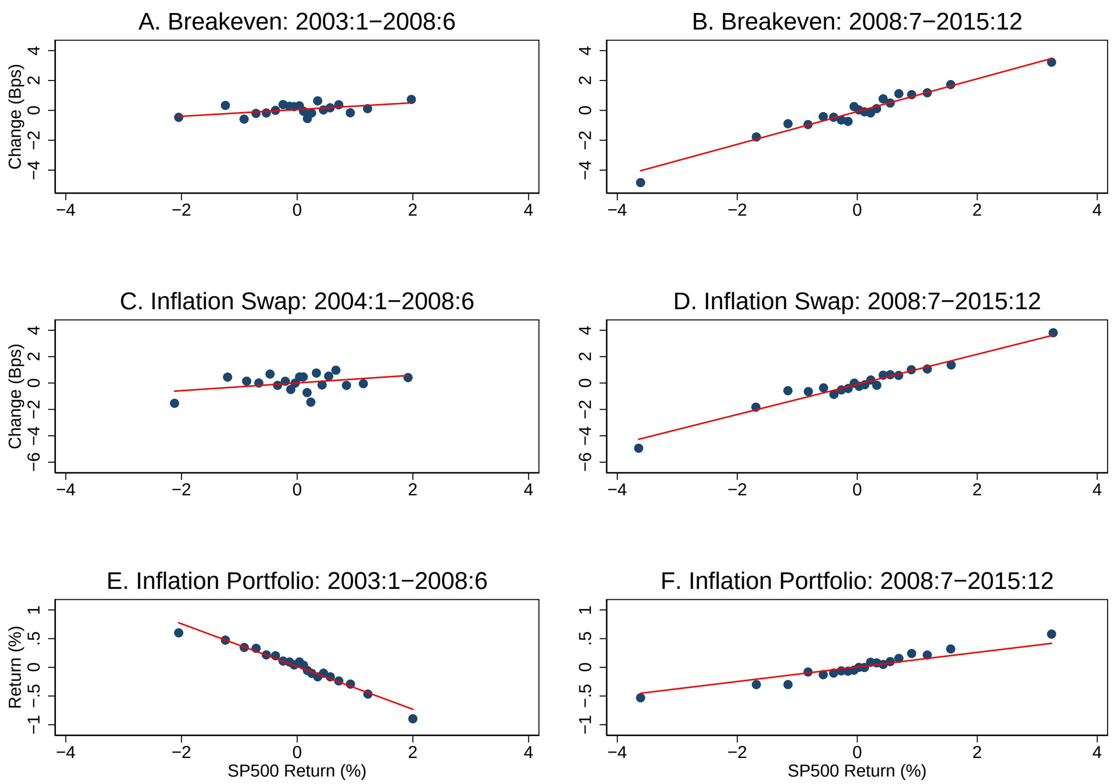

Figure 2: Binned scatter plot of daily changes in S&P 500 (x-axis) vs. daily changes in inflation compensation (y-axis) for two subsamples (left column: before July 2008; right column: July 2008 to December 2015), with regression lines superimposed. The inflation compensation is measured by 10-year breakevens in the top row, 10-year inflation swaps in the center row, and the inflation portfolio in the bottom row. See section 3.3 for details on the construction of the inflation portfolio.

2003:1-2015:12. Table 1 reports the coefficient β_2, corresponding to the association between the stock return and inflation compensation pre-2008, and the coefficient β_3, corresponding to the change after 2008, together with the associated standard errors, for a variety of inflation compensation measures: the breakeven calculated from the New York Fed's H15 release or from the Gurkaynak, Sack and Wright yield curves, from inflation swaps, or from different version of the inflation portfolio, which will be described below. The table reports results using daily data or using 20-day changes. We see that β_2 is typically small, and sometimes negative, but β_3 is always positive and highly statistically significant. Economically, a stock return of 1% is associated with an additional 0.87bps of inflation compensation. A large stock market movement, say 10%, such as those shown in Figure 1, is hence associated with a significant change in inflation compensation.

TIPS markets were disrupted during the peak of the financial crisis (see, for instance, Fleckenstein et al. (2014)). Our results hold if we exclude the period 2008:7-2009:6, as shown in Figure 1 and Table 1 in appendix, and as can be seen from Figure 1 in the introduction - the association captures cycles of good news (or optimism) and bad news (pessimism) about the recovery that are economically meaningful.

3.2 Inflation data releases

Our second piece of evidence comes from the response of the stock market to inflation data releases. We follow a large literature and regress the daily S&P 500 R_t^s return on the "surprise" component of macro announcements. The surprise, denoted $Surp_t$, is defined as the difference between the data as released by the statistical agency (here, the Bureau of Labor Statistics) and the median forecast made by economists (as collected by Action Economics/MMS). We use the same interaction regression design as in the previous section to illustrate the change in coefficient:

$$R_t^s = \beta_0 + \beta_1 D_{t\geq 2008:7} + \beta_2 Surp_t + \beta_3 D_{t\geq 2008:7} \times Surp_t + \varepsilon_t.$$

Table 2 presents the results for different inflation measures: the consumer price index (CPI), and the core CPI (which excludes food and energy prices), the production price index (PPI) and the core PPI, as well as average hourly earnings (a measure of wage inflation).[9]

[9] Because the average hourly earnings is released as part of the monthly employment report, we include as a control the employment surprise and its interaction with the post-2008 dummy. This change affects the coefficient or standard error by a very small amount.

Inflation Compensation Measure	β_2	$se(\beta_2)$	β_3	$se(\beta_3)$	N	R^2
A. Daily data						
10-year breakeven (H15)	0.23**	(0.09)	0.87***	(0.15)	3,247	0.110
10-year breakeven (GSW)	0.22**	(0.09)	0.89***	(0.15)	3,234	0.120
10-year inflation swaps	0.29*	(0.16)	0.85***	(0.23)	2,780	0.090
Inflation portfolio (ISP)	-0.37***	(0.02)	0.50***	(0.03)	3,248	0.170
Inflation portfolio, value-weighted	-0.43***	(0.02)	0.69***	(0.06)	3,248	0.110
Inflation portfolio, using CPI	-0.09***	(0.02)	0.28***	(0.03)	3,248	0.190
Inflation portfolio, excl. financials	-0.35***	(0.02)	0.46***	(0.03)	3,248	0.180
Inflation portfolio, with market control	-0.14***	(0.01)	0.25***	(0.02)	3,248	0.110
Inflation portfolio, no rebalancing	-0.39***	(0.01)	0.16***	(0.02)	3,248	0.410
B. 20-day change						
10-year breakevens (H15)	-0.36***	(0.09)	2.65***	(0.14)	3,247	0.270
10-year inflation swaps	-0.27**	(0.12)	1.73***	(0.16)	2,752	0.140
Inflation portfolio (ISP)	-0.06***	(0.02)	0.15***	(0.03)	2,542	0.030

Table 1: Changing association of inflation compensation and stock returns. The table reports the estimated coefficients from the model $\Delta IC_t = \beta_0 + \beta_1 D_{t\geq 2008:7} + \beta_2 R_t^s + \beta_3 D_{t\geq 2008:7} R_t^s + \varepsilon_t$. Each row corresponds to a different measure of inflation compensation ΔIC_t. See the text for variable descriptions. Sample is 2003:1-2015:12 except for inflation swaps, for which it is 2004:1-2015:12. White standard errors are presented in parentheses. *, **, *** denote the 10%, 5%, and 1% levels of significance.

Column 1 shows that for all inflation measures, β_2 is negative, and significantly so for all but one measure. This means that pre-2008, the stock market was negatively impacted by higher than expected inflation. Quantitatively, if CPI core inflation was one 'tick' (1/10th of a percent) higher than expected, stock prices on average fell 0.26%, a sizeable amount. Column 3 reports β_3, the estimated change in sensitivity post-2008, which is positive for all inflation measures. Given the short data span post-2008, only one measure is statistically significant. But in all cases, the estimated β_3 is so large that the sensitivity to inflation $\beta_2 + \beta_3$ actually turns *positive* after 2008.

3.3 An inflation-mimicking portfolio

We now propose an alternative measure of inflation compensation which builds on the inflation data release results. Rather than studying the response of the aggregate stock market to inflation releases, one can measure the responses of individual stocks. Some firms are naturally more sensitive to inflation, due to the nature of their business, their assets, and their liabilities.

Inflation Surprise Measure	β_2	$se(\beta_2)$	β_3	$se(\beta_3)$	N	R^2	Sample start
CPI	-0.92**	(0.39)	1.010	(1.84)	429	0.01	Feb. 1980
Core CPI	-2.67***	(0.66)	4.060	(2.69)	315	0.05	Aug. 1989
PPI	-0.28*	(0.15)	0.400	(0.33)	425	0.01	Feb. 1980
Core PPI	-0.470	(0.29)	0.750	(0.83)	304	0.01	Aug. 1989
Average hourly earnings	-0.77**	(0.38)	2.16**	(1.07)	301	0.02	Feb. 1990

Table 2: Changing association of stock returns and inflation release surprises. The table reports the estimated coefficients from the model $R_t^s = \beta_0 + \beta_1 D_{t\geq 2008:7} + \beta_2 Surp_t + \beta_3 D_{t\geq 2008:7} Surp_t + \varepsilon_t$. Each row corresponds to a different measure of inflation release surprise $Surp_t$. See the text for variable descriptions. Sample ends in 2015:12 and start date indicated in table for each surprise variable. White standard errors are presented in parentheses. *, **, *** denote the 10%, 5%, and 1% levels of significance.

We can create a long-short portfolio of stocks based on their estimated inflation sensitivity. This long-short portfolio acts as a measure of inflation compensation in the stock market.

There are two advantages to this approach over the existing inflation compensation measures: first, we can obtain a sample longer than for TIPS or inflation swaps; second this circumvents the liquidity differences that arise when constructing breakevens (e.g. nominal Treasuries are more liquid than TIPS). On the other hand, this inflation compensation measure may be more noisy as it uses stock prices rather than bond yields.

We implement this as follows. On the last day of each year, we sort the 500 stocks with largest market capitalization in CRSP by inflation sensitivity. The inflation sensitivity is estimated using the response of the stock to core CPI announcements over the previous 3 years of data. Specifically, we run for each stock the regression $R_{it} = \alpha_i + \beta_i Surp_t + \varepsilon_{it}$, over the 36 last days of CPI releases, where $Surp_t$ is the core CPI surprise. We then create an (equally-weighted) portfolio long the top quintile of inflation sensitivity and short the bottom quintile. We rebalance the portfolio every year. Financial firms and commodity producers have typically high inflation sensitivity, while tech firms tend to have low inflation sensitivity, but there is a fair amount of turnover.

The bottom panel of Figure 2 illustrates the association between the excess return on this portfolio and the S&P 500 before and after the crisis. Before the crisis, the correlation is strongly negative, i.e. this excess return has a negative CAPM market β, but it becomes strongly positive after the crisis. The change is similar to the results obtained with breakevens or inflation swaps. Table 1 presents statistical evidence for the baseline inflation portfolio and for a number of variants in the construction of the inflation portfolio: first, value-weighted rather than equal-weighted; second, using the CPI rather than core CPI to measure inflation sensitivity; third, excluding financial firms; fourth, adding the market return as a control in the measurement of inflation sensitivity, i.e. sorting firms on β_i estimated from the regression $R_{it} = \alpha_i + \beta_i Surp_t + \gamma_i R_{Mt} + \varepsilon_{it}$; fifth, not rebalancing the portfolio but keeping its composition constant after 2006. In all cases, there is strong statistical evidence of an economically meaningful change. This confirms the results obtained above with breakevens and inflation swaps using an entirely different data construction.

3.4 Association of actual inflation and stock returns

An alternative approach, that builds on a distinguished history (e.g. Fama and Schwert (1977), Modigliani and Cohn (1979)), is to directly estimate the relation between monthly (nominal) stock returns to actual inflation. Table 3 presents the results from the regression

$$R_t^s = \alpha + \beta \pi_t + \varepsilon_t,$$

for three subsamples: January 1960 through December 1983, January 1984 through June 2008, and July 2008 through December 2015. As Fama and Schwert demonstrated, higher inflation has historically been associated with lower nominal stock returns, i.e. $\beta < 0$. This stands in sharp contrast to what one might expect under the Fisher hypothesis (inflation neutrality) i.e. $\beta = 1$. The result still holds in the Great Moderation sample - an out of sample confirmation for Fama and Schwert - but disappears in the post-2008 sample, where the slope coefficient turns positive for CPI inflation.

To control for expected inflation (which may be embedded already in the stock price), the table also presents results which control for two proxies of expected inflation, the T-bill rate (used by Fama and Schwert) and one month lagged year-over-year inflation (perhaps the best simple predictor of inflation). In all cases, the results remain. Table 4 illustrates the change using the same interaction design as above, i.e.

$$R_t^s = \beta_0 + \beta_1 D_{t \geq 2008:7} + \beta_2 \pi_t + \beta_3 D_{t \geq 2008:7} \times \pi_t + \varepsilon_t,$$

estimated using either a long sample (1960-2015) that includes the Great Inflation, or using only the Great Moderation (1984-2015). In all cases, the estimated change in sensitivity is statistically significant for inflation, though not for core inflation.

Inflation and Stock Returns	1960m1-1983m12	1984m1-2008m6	2008m7-2015m12
A. Model: $R_t = \alpha + \beta\pi_t + \varepsilon_t$			
CPI	-2.84***	-3.692***	1.406
	(0.786)	(1.036)	(2.033)
Core CPI	-1.83**	-1.942	-5.98
	(0.888)	(2.162)	(7.695)
B. Model: $R_t = \alpha + \beta\pi_t + \gamma\pi_{t-13\to t-1} + \varepsilon_t$			
CPI	-5.274***	-3.784***	1.077
	(0.98)	(1.075)	(1.671)
Core CPI	-3.321***	-4.6*	-3.759
	(1.065)	(2.769)	(7.809)
C. Model: $R_t = \alpha + \beta\pi_t + \gamma Tbill_{t-1} + \varepsilon_t$			
CPI	-3.132***	-3.889***	2.247
	(0.96)	(1.054)	(1.729)
Core CPI	-1.344	-3.69	-2.24
	(1.181)	(2.395)	(7.858)

Table 3: Association of inflation and stock market returns, for different subsamples. The table reports estimates of three statistical models A, B, C, listed in the table, using either CPI or core CPI inflation, in three different subsample, using monthly data. Robust (White) standard errors in parentheses. *, **, *** denote significance at the level of 10%, 5%, and 1%.

Inflation and Stock Returns	β_2	$se(\beta_2)$	β_3	$se(\beta_3)$	N	R^2
1984-2015						
CPI	-2.407**	(1.017)	4.291**	(1.922)	384	0.024
Core CPI	-1.624	(2.067)	-2.218	(3.255)	384	0.0030
1960-2015						
CPI	-2.513***	(0.620)	4.276**	(1.911)	672	0.037
Core CPI	-1.824**	(0.806)	-2.101	(3.127)	672	0.011

Table 4: Changing association of monthly stock returns and inflation. The table reports the estimated coefficients from the model $R^s_t = \beta_0 + \beta_1 D_{t\geq 2008:7} + \beta_2\pi_t + \beta_3 D_{t\geq 2008:7}\pi_t + \varepsilon_t$, using either CPI inflation or core CPI inflation. Top panel: 1984m1-2015m12 sample; bottom panel: 1960m1-2015m12 sample. White standard errors are presented in parentheses. *, **, *** denote the 10%, 5%, and 1% levels of significance.

3.5 Correlation of inflation and economic activity

So far, we have documented substantial changes in the association of stock returns and various proxies for inflation. The standard New Keynesian model, however, speaks to the association of economic activity and inflation, or at least news about future economic activity and news about inflation. We use stock returns as a proxy for the former, because they provide a high-frequency measure of beliefs about the state of the economy. But one might prefer to use more direct measures of economic activity. We believe it is difficult to test the proposition that the association of economic activity and inflation has changed because of the short sample, where very few shocks are realized, making the signal-to-noise ratio unfavorable.[10] In contrast, in our high-frequency studies, we capture changes in beliefs about future economic activity and beliefs about future inflation, even if they do not get realized.

Nevertheless, in this section we present some simple evidence that the prevailing negative relation between inflation and economic activity turned more positive during the Great Recession. First, as a simple visual illustration, figure 3 is a scatter plot of inflation and (real, nondurables and services) consumption growth for different subsamples. Between 1959 and 1983, the relation is significantly negative. This negative relation persists during the Great Moderation period (1984m1-2008m6). It is only after July 2008 that this relation turns completely flat (or, if anything, positive).

To illustrate these changes statistically, we again use an interaction regression:

$$Y_t = \beta_0 + \beta_1 D_{t \geq 2008:7} + \beta_2 \pi_t + \beta_3 D_{t \geq 2008:7} \times \pi_t + \varepsilon_t,$$

where Y_t is a real economic activity measure, either consumption growth (as in the plots above); the growth rate of manufacturing industrial production; or the opposite of the change in the unemployment rate. Table 5 presents the results, using either total inflation and for core inflation. We also study the relation between innovations of economic activity and innovations of inflation, i.e. we replace Y_t and π_t by innovation measures.[11]

The table confirms that for many variables, the relationship was negative in the pre 2008 sample, i.e. $\beta_2 < 0$. For all but two of the twelve cases, we have $\beta_3 > 0$; four of these are statistically significant at conventional level. This suggests that there has indeed been a change in the association of economic activity and inflation.

[10] See Debortoli et al. (2019) for an attempt at assessing this change.

[11] These innovations are constructed as the residuals of a regression of economic activity (resp. inflation) on three lags of the economic activity variable and three lags of the inflation variable.

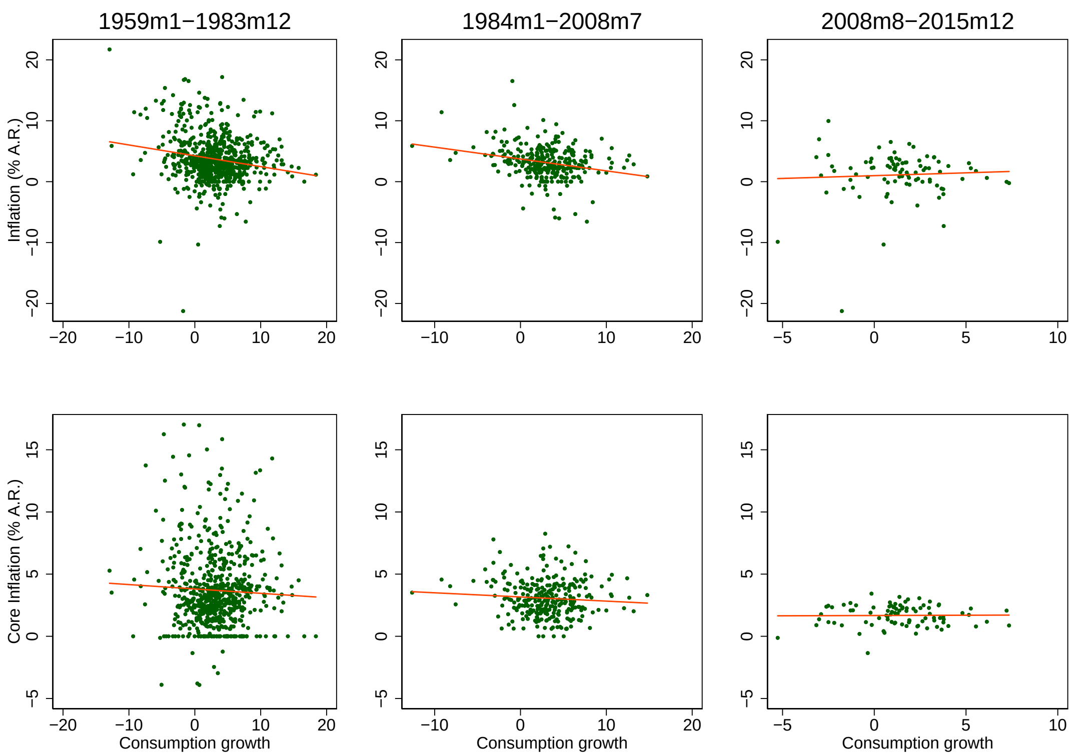

Figure 3: Association between inflation and (real nondurables and services) consumption growth. Both measures are annualized. The top panel shows the results with total inflation; the bottom panel shows the results with core inflation. Each column corresponds to a different sample period.

Economic activity and Inflation	β_2	$se(\beta_2)$	β_3	$se(\beta_3)$	N	R^2
A. CPI						
Manuf. industrial prod.	-0.301**	(0.126)	1.195***	(0.353)	672	0.0430
Consumption	-0.308***	(0.050)	0.322***	(0.097)	672	0.0960
Unemployment rate	-0.061**	(0.026)	0.224***	(0.067)	672	0.0210
B. Core CPI						
Manuf. industrial prod.	-0.275*	(0.163)	-0.279	(1.765)	672	0.0230
Consumption	-0.134**	(0.059)	0.100	(0.321)	672	0.0380
Unemployment rate	-0.0550	(0.034)	0.165	(0.314)	672	0.00600
C. CPI innovation						
Innov. manuf. industrial prod.	0.0270	(0.130)	0.0760	(0.339)	672	0.00900
Innov. consumption	-0.337***	(0.066)	0.212**	(0.095)	672	0.0660
Innov. unemployment rate	-0.00800	(0.029)	0.0390	(0.055)	672	0.00100
D. Core CPI innovation						
Innov. manuf. industrial prod.	0.584***	(0.173)	-0.583	(1.209)	672	0.0250
Innov. consumption	-0.0140	(0.085)	0.221	(0.285)	672	0.0140
Innov. unemployment rate	0.097**	(0.045)	0.0290	(0.239)	672	0.0110

Table 5: Changing association of economic activity and inflation. The table reports the estimated coefficients from the model $Y_t = \beta_0 + \beta_1 D_{t\geq 2008:7} + \beta_2 \pi_t + \beta_3 D_{t\geq 2008:7}\pi_t + \varepsilon_t$. Each panel A, B, C, D corresponds to a different inflation measure π_t, and each row within each panel corresponds to a different economic activity measure Y_t. Monthly data 1963m1 to 2015m12. White standard errors are presented in parentheses. *, **, *** denote the 10%, 5%, and 1% levels of significance.

3.6 When did a structural break occur?

The results above demonstrate that the association of inflation and stock returns or economic activity changed between the period prior to July 2008 and the period after, but we took this break date as given. In this section we provide formal statistical evidence that justifies the choice of break date. Specifically, we use the tests designed by Andrews (1993) (for a single break) and Bai and Perron (1998) (for multiple breaks), and look for breaks in the period from January 1984 to December 2015 - a standard "Great Moderation" sample during which the Fed instruments and objectives are fairly stable.

Table 6 summarizes the results of these tests.[12][13] Consider first the Andrews test, which looks for the date that generates the strongest evidence of a break. As seen in columns 2 and 3 of table 6, for seven out of the eleven associations that we study, the test selects a break date in 2008, and for one association it is in 2007. For the other three associations (surprise inflation or core inflation and stock returns, and core inflation and consumption), the test detects no statistically significant change. Perhaps more strikingly, for none of our variables does the test suggest a break in a period outside the Great Recession.[14]

The results from the Bai and Perron tests, which allow for multiple breaks, are broadly consistent with the results of the Andrews test. Regarding the association of daily stock returns and inflation, the test selects dates during the financial crisis, especially 2008, but it also finds additional changes for the inflation portfolio, e.g. 1996 and 2000, consistent with the thesis of Campbell et al. (2014). Regarding the association of monthly stock returns and inflation, the evidence is again statistically weaker, with only one break detected, in 2008. Finally, as to the association of economic activity and inflation, dates close to the financial crisis are again

[12] This footnote details how we implement the tests. The QLR test in Andrews (1993) for a single break is straightforward; we calculate the F-statistic for all potential break dates in the central 70% of the sample. The tests in Bai and Perron (1999, 2003) for multiple breaks requires more attention, especially in choosing the number of breaks. Bai and Perron (2003) show that the BIC criterion may not work well under the null of no break date, while the LWZ criterion may not work well under the alternative of some break dates. The sequential method is better than the BIC and LWZ criteria in general but it may still miss some break dates. As recommended by Bai and Perron (2003), we first use the UDmax test to see if there are any structural breaks. We then look at supF and sequential supF(.|.) test-statistics derived from the global minimization to determine the number of breaks. If the number of breaks based on the supF tests is different from the one based on the sequential method, we use the result from the supF tests. We set the maximum number of breaks $M = 5$, the trimming parameter $\varepsilon = 0.15$, the minimum length of subsample partitioned by break dates $h = T * \varepsilon$. These values are standard, and our results are quite robust to these parameters.

[13] Note that our sample for the daily breakeven and inflation swap data starts in 2003 and 2004 respectively, but our daily inflation portfolio is available since 1984. Our macroeconomic surprise data start in 1984 for inflation and 1990 for core inflation.

[14] In the online appendix, we illustrate these results graphically by showing the F-statistics for the QLR test.

	Andrews (1993)		Bai and Perron (1998, 2003)	
	QLR	Break date	UDmax	Break dates
A. Stock return				
Daily data				
Breakeven	37.63***	Feb. 1, 2008	42.35***	Aug. 22, 2008***
Inflation swap	26.4***	Aug. 7, 2007	38.96***	Dec. 7, 2007***
				Sep. 3, 2009***
				Feb. 17, 2012***
Inflation portfolio	1198.6***	Dec. 30, 2008	1380.72***	Dec. 30, 2008***
				Dec. 22, 2000***
				Jun. 25, 1996***
Monthly data releases				
Surprise inflation	6.77	Aug. 2001	5.26	No break
Surprise core inflation	10.91***	Jan. 2008	5.88	No break
Monthly data				
Inflation	13.69***	Aug. 2008	11.67**	Sep. 2008**
Core inflation	2.08	Jan. 1991	6.61	No break
B. Economic Activity				
Consumption: monthly data				
Inflation	18.38***	Jul. 2008	15.21***	Dec. 2008***
Core inflation	3.35	Jun. 2005	50.31***	Dec. 2006***
				Aug. 1990*
Industrial Production: monthly data				
Inflation	9.23**	Jul. 2008	8.08*	Jan. 1992*
				Dec. 1999*
				Sep. 2008*
Core inflation	7.66*	Dec. 2008	15.31***	Apr. 2000*
				Dec. 2007*

Table 6: Test for unknown structural breaks in the association of measures of inflation and stock returns or economic activity. The sample is January 1984 to December 2015, or whichever subset of this has data available (as noted in the text). For Andrews' QLR test, the F-statistic is computed for all potential break dates in the central 70% of the sample. For the Bai and Perron tests, we set the maximum number of breaks M=5, the trimming parameter ϵ=0.15, the minimum length of subsample partitioned by break dates $h = T \times \epsilon$. *, **, *** denote the 10%, 5%, and 1% levels of significance.

selected, and additional breaks are found in 1990 or 1992, and 1999 or 2000.

In the online appendix, we illustrate graphically the changes in association using rolling regressions, several of which exhibit large swings in coefficients around 2008. We also depict the QLR statistic and report the results using a longer sample that starts in January 1959 for the monthly data only.

Overall, the evidence is strong that there was a break in many of these associations around 2008, though there is some uncertainty about the exact date. There is also significant evidence that additional breaks occurred in the 1990s, presumably reflecting secular changes such as those studied by Campbell et al. (2014).

4 Model

Our model builds on the standard New Keynesian model as outlined for instance in Gali (2008) and Woodford (2003). Like Rudebusch and Swanson (2012), we use recursive preferences (Epstein and Zin (1989)) with high risk aversion so that the model replicates some basic properties of the yield curve. Compared to that paper, the main difference is that we explicitly take into account the zero lower bound.[15] In the interest of transparency, and given the difficulty in solving the model (numerically), we abstract from capital accumulation, and from wage stickiness and other frictions, even though they are no doubt important in reality.

4.1 Household

The representative household works, consume, and decides how much to save in various assets. These assets are in zero net supply (since there is no capital). We follow Rudebusch and Swanson's version of recursive preferences. Denoting by V_t the intertemporal utility and by $u(C_t, N_t)$ the flow utility of consumption C_t and labor N_t, we assume that:

$$V_t = (1-\beta)\, u(C_t, N_t) + \beta E_t \left(V_{t+1}^{1-\alpha}\right)^{\frac{1}{1-\alpha}},$$

and

$$u\left(C_t, N_t\right) = \frac{C_t^{1-\sigma}}{1-\sigma} - \frac{\chi N_t^{1+\nu}}{1+\nu},$$

If the parameters we use lead to a negative flow utility $u(C_t, N_t)$, we define utility as:

$$V_t = (1-\beta)\, u(C_t, N_t) - \beta E_t \left(\left(-V_{t+1}\right)^{1-\alpha}\right)^{\frac{1}{1-\alpha}}.$$

[15] The other differences are that we use Rotemberg rather than Calvo pricing to economize on state variables, and we also use different shocks and a different monetary policy rule.

The household budget constraint is:

$$P_t C_t + \xi_t B_t + Q_t S_t = W_t N_t + \Pi_t + R_{t-1} B_{t-1} + (Q_t + D_t) S_{t-1},$$

where S_t is the number of shares purchased at time t, Q_t is the stock price, D_t is the dividend per share, B_t is the quantity of one-period risk-free assets, Π_t are firms' profits, rebated to the household, and W_t is the wage rate.

We make two simplifying assumptions. First, we assume that the shares are not claims to the firms' profits but, as in much of the asset pricing literature, a levered claim on consumption, $D_t = C_t^{\psi}$, where $\psi \geq 1$ is a parameter capturing financial and operating leverage (Abel (1999)). This assumption circumvents a well-known deficiency of the New Keynesian model: firm profits are much less cyclical than in the data.[16] Second, we assume that the household purchases risk-free assets at a discount ξ_t, which is an exogenous stochastic process. We interpret this discount as reflecting the convenience yield of safe and liquid assets, which has been emphasized as an important factor in recent research (for instance, Krishnamurthy and Vissing-Jorgensen (2012)). This convenience yield could be motivated by introducing liquidity in the utility function as in Fischer (2014), but for simplicity we introduce it as a direct subsidy to risk-free assets. This shock will plays the same role in our model as the preference shock used in much of the New Keynesian literature. We will refer to ξ_t as "demand shock" or "liquidity shock" equivalently. Importantly, we will assume that this convenience yield applies to all risk-free assets, regardless of their maturity, and whether they are nominal or real (inflation-indexed).

The labor supply equation is simply

$$W_t = \frac{u_2(C_t, N_t)}{u_1(C_t, N_t)} = \chi C_t^{\sigma} N_t^{\upsilon}. \tag{1}$$

The real stochastic discount factor is

$$M_{t+1}^r = \beta \left(\frac{C_{t+1}}{C_t} \right)^{-\sigma} \left(\frac{V_{t+1}}{E_t \left(V_{t+1}^{1-\alpha} \right)^{\frac{1}{1-\alpha}}} \right)^{-\alpha},$$

and the nominal stochastic discount factor is

$$M_{t+1}^n = \frac{M_{t+1}^r}{\Pi_{t+1}},$$

where Π_{t+1} is gross inflation P_{t+1}/P_t.

[16] A number of extensions have been proposed to explain this cyclicality, for instance fixed costs, sticky wages, or financial leverage (see, for instance, Li and Palomino (2014)). We do not incorporate these extensions in the interest of simplicity.

The first order condition links the nominal short-term interest rate to the nominal stochastic discount factor:

$$1 = E_t \left[\xi_t^{-1} R_t M_{t+1}^n \right],$$

where $R_t \equiv Y_t^{n,(1)}$ is the gross nominal yield on a 1-period risk-free bond.

4.2 Production and price-setting

Our modeling of the production side is completely standard from the New Keynesian literature. There is a measure one of identical monopolistically competitive firms, each of which operates a constant return to scale, labor-only production function:

$$Y_{it} = Z_t N_{it}, \tag{2}$$

where Z_t is an exogenous stochastic productivity process, common to all firms. Each firm faces a downward-sloping demand curve coming from the Dixit Stiglitz aggregator with elasticity of demand ε:

$$Y_{it} = Y_t \left(\frac{P_{it}}{P_t} \right)^{-\varepsilon}, \tag{3}$$

where P_t is the price aggregator:

$$P_t = \left(\int_0^1 P_{it}^{1-\varepsilon} di \right)^{\frac{1}{1-\varepsilon}}.$$

We use the Rotemberg (1982) assumption of quadratic adjustment costs to changing prices.[17] Specifically, the cost of changing the price from P to P' is $\frac{\phi}{2} Y \left(\frac{P'}{P} - \overline{\Pi} \right)^2$ where ϕ captures the magnitude of the costs, Y are firm sales, and $\overline{\Pi}$ is a parameter capturing "inflation indexation".

Each period, firms set their price P_{it} so as to maximize

$$E_t \sum_{k=0}^{\infty} M_{t,t+k}^n \left(P_{it+k} Y_{it+k} - W_{t+k} N_{it+k} - \frac{\phi}{2} Y_{it+k} \left(\frac{P_{it+k}}{P_{it+k-1}} - \overline{\Pi} \right)^2 \right),$$

subject to the the production function (2) and the demand curve (3).

In equilibrium, all firms choose the same price, and given quadratic adjustment costs, they adjust their price each period. Taking first order conditions yields the standard Rotemberg forward-looking Phillips curve:

$$0 = \left(1 - \varepsilon + \varepsilon \frac{W_t}{Z_t} - \phi(\Pi_t - \overline{\Pi})\Pi_t \right) Y_t + \phi E_t \left(M_{t+1}^r (\Pi_{t+1} - \overline{\Pi}) \Pi_{t+1} Y_{t+1} \right).$$

[17] Miao and Ngo (2014) illustrate that some results (such as the size of the fiscal multiplier) may be affected by the price setting assumptions at the ZLB.

Finally, the resource constraint reads

$$C_t = \left(1 - \frac{\phi}{2}(\Pi_t - \overline{\Pi})^2\right) Y_t, \tag{4}$$

since we need to subtract price adjustment costs from output. Measured gross domestic product equals consumption, since price adjustment is an intermediate input: $GDP_t = C_t$.

4.3 Fundamental shocks

We assume that both the liquidity and the productivity shock follow independent AR(1) processes with normal innovations:

$$\log \xi_t = \rho_\xi \log \xi_{t-1} + \varepsilon_{\xi,t},$$

with $\varepsilon_{\xi,t}$ i.i.d $N(0, \sigma_\xi^2)$, and

$$\log Z_t = \rho_z \log Z_{t-1} + \varepsilon_{z,t},$$

with $\varepsilon_{z,t}$ i.i.d $N(0, \sigma_z^2)$. Throughout the paper we use "liquidity shock" or "preference shock" or "demand shock" equivalently to refer to $\varepsilon_{\xi,t}$, and "supply shock" or "productivity shock" to refer to $\varepsilon_{Z,t}$.

4.4 Monetary policy

We assume that monetary policy follows the following rule, in the spirit of Taylor (1993):

$$R_t = \max\left\{1, R^* \left(\frac{\Pi_t}{\Pi^*}\right)^{\phi_\pi} \left(\frac{GDP_t}{GDP^*}\right)^{\phi_y}\right\} \tag{5}$$

where R_t is the gross nominal interest on a one-period risk-free bond; ϕ_π and ϕ_y are the responsiveness to inflation and GDP respectively; and R^*, Π^* and GDP^* are constants.[18] The max operator reflects the truncation implied by the zero lower bound.

Taylor's original rule (1993) assumes that the central bank responds to the deviation of GDP from potential GDP, i.e. the level of GDP that would prevail in an economy without price stickiness; in contrast, we assume (as in Fernandez-Villaverde et al. (2015) and Swanson (2015a)) that the central bank responds to the deviation of GDP from "trend".[19] This can be

[18] We set Π^* and GDP^* equal to the nonstochastic steady-state values of inflation and GDP, and adjust R^* so that the model generates an average inflation equal to 2%.

[19] Given that our model abstracts from long-run growth, actual GDP and potential GDP are both stationary. Trend GDP is simply GDP^*. Potential GDP is affected by current productivity; it can be shown that potential GDP equals $GDP^* Z_t^{\frac{1+\upsilon}{\sigma+\upsilon}}$.

motivated by the difficulty of measuring potential GDP, in particular in real time. We discuss in section 5.5 the importance of this assumption for our results.

We also abstract from so-called "unconventional" policies such as forward guidance or asset purchases (QE or LSAP). There is a significant debate on the efficacy of these policies. It is usually believed that their effects are less potent or more uncertain, than those of traditional interest rate policy. On top of this, because of the political risk, central banks may be reluctant to use them fully.[20]

4.5 Asset prices

This section describes the various asset prices that we calculate in the model.[21] In order to simplify the numerical computation of the model, we study the prices of geometric consols rather than zero-coupon bonds. A nominal geometric consol with parameter λ pays \$1 next period, \$$\lambda$ the period after, \$$\lambda^2$ the period after that, and so on. A real consol with parameter λ has the same payoffs, but in units of final goods rather than in dollars. We can then choose λ to obtain an asset with the same (average) duration as a reference asset we want to compare it to in the data, for instance the 10-year Treasury note.

The geometric nature of payoffs implies that the consol price satisfies the recursion:

$$q_t^{i,\lambda} = E_t \left[\frac{1}{\xi_t} M_{t+1}^i \left(1 + \lambda q_{t+1}^{i,\lambda} \right) \right] \tag{6}$$

where we denote by i the nominal or real nature of the consol, $i \in \{n, r\}$, and note that the consols are also subject to the liquidity convenience ξ_t. The holding period (gross) return on a consol is given by:

$$R_{t+1}^{i,\lambda} = \frac{1 + \lambda q_{t+1}^{i,\lambda}}{q_t^{i,\lambda}}, \tag{7}$$

and the (gross) yield is defined as

$$Y_t^{i,\lambda} = \frac{1}{q_t^{i,\lambda}} + \lambda. \tag{8}$$

Inflation breakevens are the difference between the log nominal yield and the log real yield

$$BE_t^{\lambda} = \log Y_t^{n,\lambda} - \log Y_t^{r,\lambda}. \tag{9}$$

[20] We also abstract from fiscal policy, which may play an important role to stimulate output at the zero lower bound.

[21] The household budget constraint includes only short-term bonds and stocks. However, we can introduce any asset in zero-net supply to calculate its price, because the introduction of the asset has no effect on the equilibrium of the model.

Next, we calculate the risk-neutral price, i.e. the price that would occur if agents priced the asset as if they were risk-neutral:

$$q_t^{i,\lambda,RN} = E_t\left[\left(1+\lambda q_{t+1}^{i,\lambda,RN}\right)\right] E_t\left[\frac{1}{\xi_t} M_{t+1}^i\right], \tag{10}$$

and from this price we can define the risk-neutral yield

$$Y_t^{i,\lambda,RN} = \frac{1}{q_t^{i,\lambda,RN}} + \lambda.$$

We define *nominal and real term premia* as the difference between the log yield and the log risk-neutral yield, respectively for the nominal and the real consol:

$$TP_t^{i,\lambda} = \log Y_t^{i,\lambda} - \log Y_t^{i,\lambda,RN}, \tag{11}$$

where $i \in \{n, r\}$. We define the *inflation term premium* as the difference between the nominal term premium and the real term premium:

$$ITP_t^{\lambda} = TP_t^{n,\lambda} - TP_t^{r,\lambda}. \tag{12}$$

Turning to stocks, as explained in section 4.1, we define, following Abel (1999), a stock as an asset with payoff $D_t = C_t^{\psi}$, where $\psi \geq 1$ reflects leverage. The real stock price satisfies the usual recursion $P_t^s = E_t\left[M_{t+1}\left(P_{t+1}^s + D_{t+1}\right)\right]$, and stock returns as $R_{t+1}^s = (P_{t+1}^s + D_{t+1})/P_t^s$.

5 Quantitative results

This section studies the quantitative implications of the model presented in the previous section. We first discuss our choice of parameters. We then explain the key economic mechanisms by showing how the response of the economy to either supply (productivity) or demand (liquidity) disturbances changes when the ZLB binds. We next show how, as a result of these changing responses, the model generates a significant change in risk premia as the economy approaches the ZLB. We then evaluate quantitatively the performance of the model, and illustrate how our results depend on various features.

5.1 Parametrization and solution method

Table 7 presents the baseline parameters that we use for our quantitative analysis; section 5.5 provides comparative statics. Most of the parameters are taken from the New Keynesian literature. The time period is one quarter. The time discount factor β is 0.992, in line with

Parameter	Description and source	Value
β	Subjective discount factor	0.992
α	Curvature with respect to next period value (note: CRRA=136)	-190.00
σ	IES is 0.5	2.00
υ	Frisch labor supply elasticity is 0.66	1.50
χ	Calibrated to achieve the steady state labor of 1/3	40.66
ε	Gross markup is 1.15	7.66
ϕ_π	Weight on inflation in the Taylor rule	2.00
ϕ_y	Weight on output in the Taylor rule	0.13
Π^*	Inflation target	1.020
R^*	Taylor rule intercept	1.051
ϕ	Adjustment cost, corresponding to the Calvo parameter of 0.85	238.11
ρ_z	Persistence of technology shock	0.92
ρ_ξ	Persistence of demand shock	0.90
σ_z	Std. dev. of the technology innovations (%)	0.63
σ_ξ	Std. dev. of the preference innovations (%)	0.14

Table 7: Model parameters.

Woodford (2003, 2011), generating an average annualized real interest rate of 2.9%.[22] We set the intertemporal elasticity of substitution (IES) of consumption $1/\sigma$ to 0.5, and the Frisch elasticity of labor supply to 2/3, again in line with the literature. We also set the gross markup to 1.15, corresponding to the the demand elasticity parameter $\varepsilon = 7.66$. The adjustment cost of changing prices is $\phi = 238$, which maps in a Calvo model to a probability of keeping price unchanged of 0.85 per quarter. This value is a little above observed price stickiness in the data, but is consistent with the estimates of Del Negro et al. (2015) and lower than those of Leeper et al. (2017).

Our monetary policy rule builds on recent empirical estimates by Gust et al. (2017) and Arouba et al. (2018); the weight on inflation in the Taylor rule is $\phi_\pi = 2$ and the weight on output gap is ϕ_y equal to 0.13 (which translates into the usual 0.5 response once the interest rate is annualized).

We set Π^* (the so-called target inflation rate) to the conventional value of 2% ($\Pi^* = 1.02$) and adjust the intercept of the Taylor rule, R^* to match the average inflation rate in the model to 2%, leading to $R^* = 1.051$. We also set $\overline{\Pi} = \Pi^*$ as is standard.

Finally, the shock process we choose is also in line with the New Keynesian literature: both productivity and liquidity shocks play an important role and both are fairly transitory

[22] Note that this average is less than the nonstochastic steady state value of 3.3% due to precautionary savings.

(in contrast to the asset pricing literature which largely focuses on shocks that have very persistent effects). The persistence of technology shocks is 0.92 and the standard deviation of the innovation is 0.63%, in line with Anzoategui et al. (2017). The persistence of the liquidity shock is 0.90 and the standard deviation of the innovation is 0.14%, in line with Gust et al. (2017) and Arouba et al. (2018). Overall, these shock processes (together with the other model parameters) imply an unconditional probability of hitting the ZLB around 5% (see Fernandez-Villaverde et al. (2015)).

We diverge strongly from the New Keynesian consensus on one parameter: risk aversion, which is related to the parameter α which measures the curvature with respect to next period value in the recursive preference. Given our shock process, consumption volatility is fairly low. As a result, the model requires a high risk aversion to generate sizeable risk premia. We set α to -190, which corresponds to a relative risk aversion to consumption (CRRA) of 137 once we take into account the curvature parameters on consumption and labor in the flow utility (see Rudebusch and Swanson (2012), Swanson (2015b)). Clearly, this parameter does not reflect the preferences of any single individual. Rather, it captures the aversion of the macroeconomy to fairly small fluctuations in aggregate consumption, as inferred from asset prices. The value we use is actually relatively modest; for instance Swanson (2015b) requires α to be -338, or CRRA to be 600, to generate the equity premium of only 1.5% per annum. Rudebusch and Swanson (2012) require α to be -396, or CRRA to be 200, to generate a term premium of 1.06% in line with the U.S. data.[23] We will discuss below in detail how this risk aversion affects macroeconomic dynamics and other model features.

Due to the presence of the ZLB, we need to solve the model using nonlinear methods. This is especially important because asset prices can be sensitive to nonlinearities. We use projection methods with cubic spline, similar to Fernandez-Villaverde et al. (2015), Miao and Ngo (2014), and Ngo (2018). Our solution method is detailed in the online appendix.

5.2 Time-varying response to macroeconomic shocks

This section discusses the effects of the two fundamental drivers of our model - technology and liquidity shocks, and in particular how the responses differ depending on the initial condition of the economy. We start by discussing the responses when the economy starts in the nonstochastic

[23] From Swanson (2015b), to generate the equity premium of 1.5%, which is smaller than the value estimated by US data, α and CRRA are required to be around -338 and 600, respectively, based on figure 1. According to Rudebusch and Swanson (2012), to generate the value of term premia of 1.06, which is in line with US data, the CRRA is required to be approximately 200 based on figure 1. This means α is approximately -396, based on the formula in footnote 23 of this paper. Note that the IES and Frisch values in our paper are the same as those in their paper.

steady-state (i.e., far from the zero lower bound) before discussing the responses when the economy is already at the zero lower bound.

5.2.1 Responses far from the Zero Lower Bound

Productivity shock

As shown by the solid black lines in the right panels in figure 4, an increase in productivity leads to higher consumption (which equals GDP), and hence lower marginal utility. The stock price increases, since it equals a present discounted value of future dividends, which are proportional to consumption. Inflation falls, since real marginal costs fall due to the higher productivity - a pattern typical of the New Keynesian model. To put it another way, price rigidities prevent a full expansion of output to its new higher potential, leading to low inflation. As a result, the covariance of consumption growth and inflation is negative, generating a positive inflation risk premium. This is the well-known "supply shock" view of inflation: inflation is driven by low productivity, and hence is countercyclical.

The effect of productivity on interest rates depends both on the monetary policy rule and on the time series process for the productivity shock. Given that the later is mean-reverting, real interest rates tend to go down when the level of productivity is high, since agents rationally expect lower real consumption growth in the future, leading to higher desired savings today. The nominal rate (or policy rate) hence falls more than inflation. This explains why in figure 4 the policy rate falls in response to the productivity shock. Another way to understand the same result is that the decline of the policy rate stimulates demand, which is necessary to bring it in line with the newly expanded supply.

Because the shock is persistent, the yields on long-term nominal and real bonds fall as both inflation and real rates fall persistently. These decline in the yields generate a positive excess return on long-term (nominal or real) bonds on impact, i.e. when the productivity shock hits. Because bonds have positive excess returns in times when marginal utility of consumption is low, they are risky assets, and hence will earn a positive risk premium, which means that the yield curve is upward-sloping for both nominal and real bonds.

Liquidity shock

Let us turn now to the liquidity shock. As shown by the solid black lines of the left panels of figure 4, an increase in the demand for liquid assets leads to decline of inflation, consumption and stock prices. The mechanism is standard: given higher demand for liquid assets, agents want to save more. The central bank reduces the short-term interest rate, to meet this increased demand. But this reduction is too small given our (standard) monetary policy rule, and the

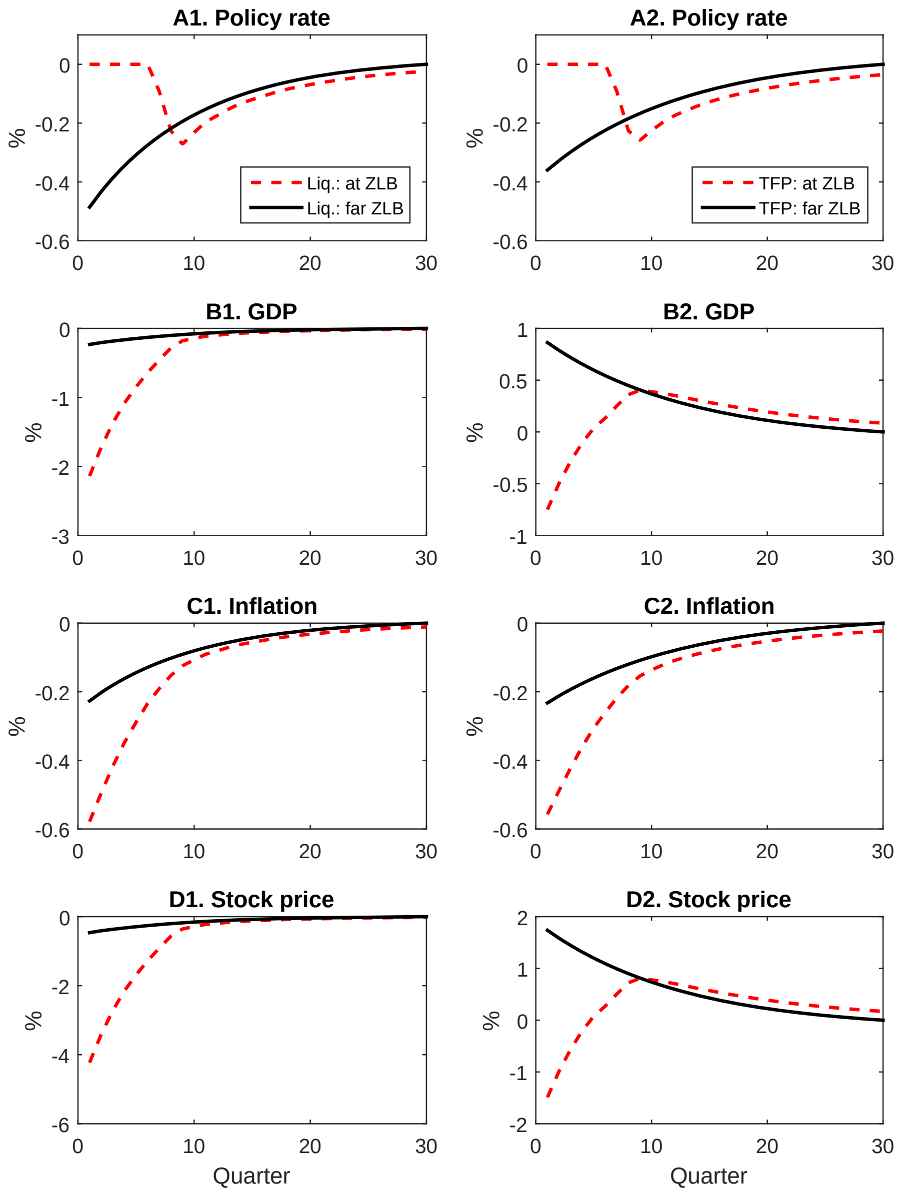

Figure 4: Impulse response to a one standard deviation liquidity shock (left column) or productivity shock (right column) when the economy is at the ZLB (red dashed line) vs. in steady-state (i.e., far from the ZLB, black full line).The dashed red lines present impulse responses at the ZLB, which is calculated as the difference between two parths: (i) a path with only large liquidity shock that brings the economy to the ZLB, and (ii) a path with the same shock, plus an additional one-standard-deviation shock.

higher desire for savings (i.e. relatively lower demand for goods) translates into lower output given sticky prices. Inflation and consumption are now positively correlated, so the inflation risk premium is negative. This is the "demand view" of inflation, which is caused by demand and hence is procyclical.

The real interest rate rises since consumption is expected to recover, but the nominal interest rate falls as lower expected inflation dominates the effect of the real rate. Overall, long-term nominal yields fall, hence long-term nominal bonds have a positive excess return on impact of a liquidity shock, which implies they will earn a negative risk premium since they provide good returns in "bad times" when marginal utility is high.

Overall, if there were only liquidity shocks, the inflation premium and nominal term premium would be negative because inflation is low when marginal utility is high (consumption is low). If there were only productivity shocks, they would be positive. For our calibration, the productivity shock is on average more important, leading to a positive inflation premium in normal times. This is consistent with an upward-sloping nominal yield curve.[24]

5.2.2 Responses at the Zero Lower Bound

We now illustrate that the responses to the same macroeconomic shocks can be quite different if the economy has a different initial condition. To do so, we calculate the response to the same shock as above, but when the economy is already at the ZLB, and compare it to the steady-state response discussed in the previous section. The impulse response at the ZLB is calculated as the difference between two paths: (i) a path with a large liquidity shock that brings the economy to the ZLB, and (ii) a path with the same shock, plus a one-standard deviation shock to either liquidity or technology. The difference gives us the effect of a one-standard deviation shock at the ZLB.[25]

Liquidity shock

The response to a liquidity shock if the ZLB binds is shown as the dashed red lines in the left panels of figure 4. Note the large difference with the black line, which is the response to the same shock if the economy starts far from the ZLB (at steady state). The most obvious difference

[24] As for the real yield curve, note that the demand shock generates an upward-sloping real yield curve, because long-term real yields rise in response to the demand shock. The real term premium is hence upward-sloping owing both to productivity and demand shocks. This does not mean, however, that the real yield curve slopes up more than the nominal yield curve in the model, because the productivity shock generates a more strongly upward-sloping curve for nominal bonds than for real bonds.

[25] Note that this is not a generalized impulse response function (GIRF) which would average over the realizations of potential future shocks, but rather a particular (example) of sample path.

is that the short-term nominal interest rate (or policy rate) cannot respond at the ZLB. As a result, there is no decline in the interest rate to offset the higher desire for savings in safe assets. This leads consumption and inflation to drop much more significantly than if the ZLB was slack. Stock prices fall since dividends (assumed to be proportional to consumption) fall. Clearly, the covariance of stock returns and inflation (or consumption and inflation) implied by this shock is much larger at the ZLB.[26]

Productivity shock

Similarly, the dashed red lines of the right panels in figure 4 displays the effect of a one-standard-deviation productivity shock when the economy starts at the ZLB. In normal times, higher TFP leads to higher consumption, lower inflation, and a lower nominal interest rate. However, if the ZLB binds, the interest rate cannot respond. Hence, inflation will fall much more as demand cannot expand to meet the new supply. This lower inflation in turn increases the real interest rate, leading consumption to actually fall on impact. This result is sensitive to parameters; what is generic is that consumption rises less than if the ZLB were not binding. In general, the overall effect of the ZLB on the covariance of consumption (or stock returns) and inflation is uncertain, but it tends to increase (become less negative, or even positive) because of the lower increase in consumption initially.[27]

Overall, the effect of the ZLB is to increase the covariance of stock returns and inflation (or consumption and inflation). We next show this point directly and draw out the implications for asset prices.

5.3 Changes in covariances and risk premia

Figure 5 visually illustrates the state-dependence in covariance and risk premia. The two axis are the two state variables of the model, which are both exogenous: the current value of productivity Z_t and of liquidity ξ_t. Panel A depicts the current policy rate; the ZLB binds when

[26] One point that might confuse the reader is that the effect of the shock is to make the interest rate fall in period 8, even though the ZLB is binding. This is because we study the effect of an additional shock in an economy that is at the ZLB in period 0, by showing the difference between two economies that both start at the ZLB, one of which faces an additional shock. Both economies over time revert to positive interest rates, but the one that has the additional shock naturally "lifts off" later. The decline in the interest rate, then, corresponds to the time period when the economy without the additional shock lifts off, while the other one remains at the ZLB. This corresponds to the first period where policy is actually able to react to the additional shock.

[27] There is a large macroeconomics literature debating the empirical relevance of the New Keynesian model dynamics (see in particular Wieland (2014)). However, it is important to note that for the purpose of this paper, we do not actually require that consumption falls with positive productivity shocks at the ZLB. It is enough that consumption increases less, and inflation decreases more, to affect the key covariance of consumption and inflation. It is even enough for us that the demand shock propagation changes, without any change to the propagation of supply shocks.

the economy is in the Northeast quadrant. Panel B depicts the long-term nominal interest rate (corresponding to a geometric consol with the maturity equivalent of a 10-year Treasury note), which moves in a similar fashion to the policy rate, though by smaller amounts.

Panel E depicts the conditional covariance of consumption growth and inflation, which is negative in normal times when the ZLB is not binding, but rises substantially when the economy operates close to the ZLB. The conditional covariance between stock returns and breakevens also tends to rise when the economy becomes closer to the ZLB, as seen in Panel F, and so does the (not depicted) covariance of stock returns and inflation. This suggests that the model is, at least qualitatively, consistent with the empirical findings of Section 3. (We discuss the quantitative fit in the next section.)

Panels C and D of Figure 5 depict the nominal term premium and the inflation premium for a 10-year equivalent geometric consol. These premia are positive in normal times, but becomes smaller when the economy is close to the ZLB. This reflects the large change in the conditional covariance of consumption and inflation together with the high risk aversion. The inflation term premium even becomes negative when the economy is deep into the ZLB territory.

To understand the change in nominal and real bond premia, recall that the TFP shock generates a positive bond premium while the liquidity shock generates a negative one. At the ZLB, consumption (and hence marginal utility) reacts much more to the liquidity shock, which increases the magnitude of the demand-shock induced risk premium - making it more negative. Conversely, consumption becomes less sensitive to TFP, which reduces the positive term premium from the TFP shock. On top of that, inflation becomes more procyclical as discussed above. Overall, these effects tend to reduce bond premia. Put another way, bonds become a better hedge for macroeconomic shocks, and hence their prices rise, i.e. the risk premium falls.

To further understand the model mechanics, figure 6 presents an illustrative simulation where the ZLB binds for 25 periods starting from period 30. The ZLB binds due to a series of liquidity shocks - i.e., the demand for "safe assets" increases. The nominal short- and long-term (10-year) yields, the nominal long-term premia, and the inflation term premium decline when the economy approaches the ZLB, then stay relatively low during the ZLB.

5.4 Quantitative evaluation of the model

The previous section shows that the model generates qualitatively a changing association of stock returns and inflation, and that in turn this implies a change in risk premia. In this section we evaluate the quantitative magnitude of these changes. We simulate the model for 500,000

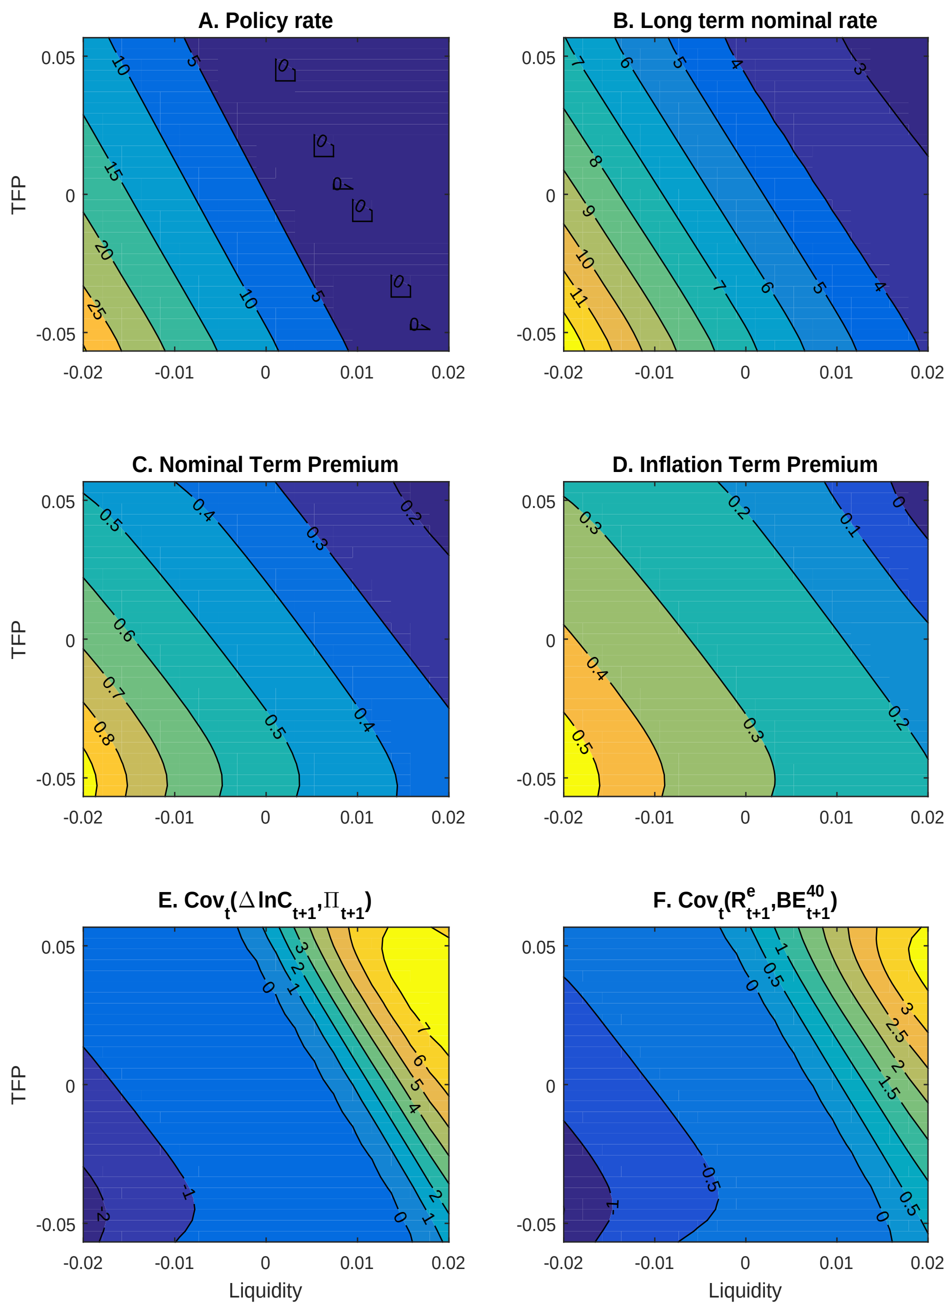

Figure 5: Term premia and conditional covariances as a function of TFP and preference shocks.

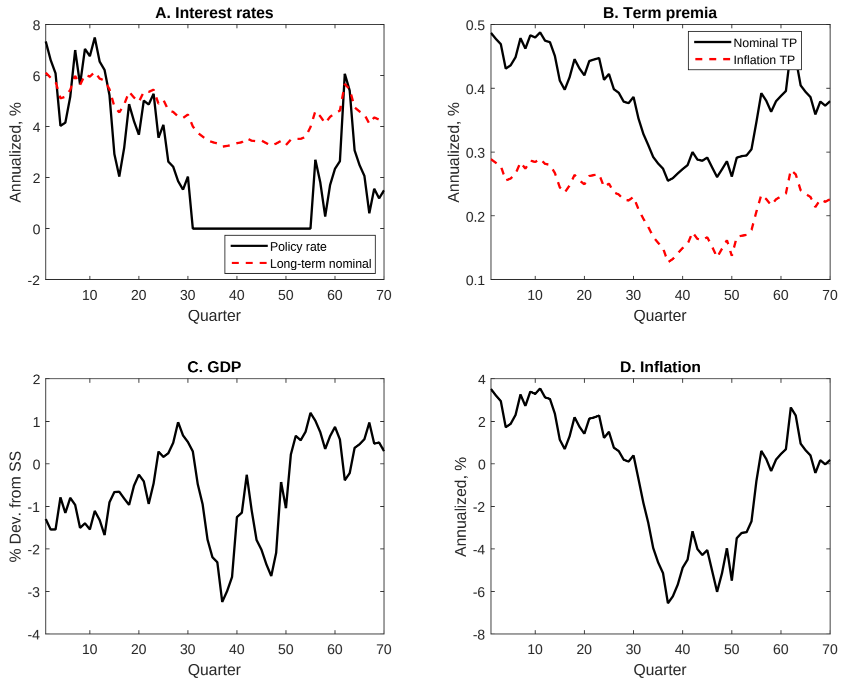

Figure 6: Simulated paths for selected economic variables. In this example, the ZLB binds for 25 periods starting in period 30.

periods (quarters) and report in Table 8 some moments of interest. Because our interest lies in how the behavior of the model changes as the economy approaches and hits the zero lower bound, we report the moments both for the full sample and for subsamples which differ on how far from the ZLB the economy is. We define the subsamples based on the implied Taylor (1993) interest rate rule:[28] less than -2%, or over 4%. We use the Taylor interest rate because it is a simple measure of how binding the ZLB is: at the depth of the Great Recession, the Taylor rule implied a policy rate around -5%.

Consider first the full sample moments. The model generates a reasonable volatility for the output gap and employment - the output gap is slightly more volatile than the data, while labor is slightly less volatile. Average (annualized) inflation is 2%, which follows from our choice of R^*, and is close the the data (1.8%). The volatility of inflation is roughly in line with the data (1.5% vs. 1.2%). The average yield on the 10-year Treasury note in our model matches the data (5.5%), while our short-term rate is higher than the data. Hence, while our model produces an upward sloping yield curve on average (37bps), it is flatter than in the data

[28] In our model, the simple Taylor rule is: $R_t^{TR} = 2 + 1.5(\pi_t - \pi^*) + 0.5 GDP_t/(GDP^* \times Z_t^e)$, since GDP potential is $GDP^* * Z_t^e$.

(184bps).[29] The nominal term premium is on average 44bps,[30] and the real term premium is lower at 18bps, so our model implies an inflation term premium of 26bps. Hence, the average breakeven, at 2.23% per year, is higher than average inflation, reflecting that *on average* agents fear inflation, because it occurs in states with high marginal utility.

Turn now to the changing association of inflation and stock returns, which we documented in Section 3. To evaluate the ability of our model to capture this fact, we run in our model the same regressions we estimated in Section 3 in actual data. The results are presented in table 9, which show the slope coefficient in a regression of stock returns on the change in breakeven, for each subsample of the model, and for the data. In the "deep ZLB" subsample of the model, where the Taylor rate is below -2%, this coefficient is about 6.8. Once the economy is far from the ZLB, this coefficient drops to -2.2. These coefficients are broadly similar to what we observe in our pre-crisis and post-crisis sample: 9.6 and 2.0 using inflation swaps, 12.7 and 2.6 using breakevens. Hence, our model replicates reasonably well this changing association.

How does this changing association in turn affect asset prices? Return to Table 8 and consider the "deep" zero lower bound subsample, when the Taylor rate is below -2%. The nominal and inflation premia are smaller in that sample, by about 16bps and 11bps, respectively, than in the full sample, and smaller by 29bps and 19bps, respectively, than in the sample "far from the ZLB".[31] By contrast, a standard New Keynesian model with low risk aversion would imply that these term premia are constant and essentially nil.

Hence, term premia move meaningfully as the economy gets closer to the ZLB, and this decline contributes to the decline in the long-term interest rate and in the breakeven. However, this contribution is quantitatively limited. The long-term nominal interest rate is almost 200bps smaller in the "deep ZLB" subsample than in the full sample, which dwarfs the 16bps mentioned above. To put it another way, at most 10% of the decline in interest rates is attributed to the lower term premium - rather than expected inflation or expected real interest rates.

We also need to highlight some limitations of the model. The ZLB period is associated with low output and employment, as is standard in New Keynesian models: for instance, the mean output gap is -2.2% in the ZLB subsample. Indeed, the model implies a severe deflation

[29] One reason why the average slope is too low in the model is that we abstract from a special liquidity preference for Treasury bills.

[30] This number is not exactly equal to the average slope because we report the average slope as the difference between average log gros yields, and similarly the term premium is the average difference between the log gross yield and the gross risk-neutral yield.

[31] The real term premium also compresses, owing to the changing consumption dynamics at the zero lower bound.

	Data		Full sample		$i^{TR} \leq -2\%$	$i^{TR} \geq 4\%$
	Mean	Sd	Mean	Sd	Mean	Mean
logOutputGap	0.00	1.06	0.03	1.50	-2.25	7.00
logLaborGap	0.00	1.14	0.01	0.75	-1.61	2.80
π	1.80	1.20	2.01	1.50	-4.53	6.18
$y^{n(1)}$	3.60	2.60	4.98	2.57	0.00	12.14
$y^{n(40)}$	5.50	2.30	5.35	0.77	3.39	7.58
$y^{r(1)}$	NaN	NaN	2.94	1.28	3.73	6.34
$y^{r(40)}$	3.30	0.90	3.12	0.32	2.66	4.01
$BE^{(40)}$	2.20	0.40	2.23	0.46	0.72	3.57
Real term premium	NaN	NaN	0.18	0.02	0.12	0.23
Nominal term premium	NaN	NaN	0.44	0.05	0.28	0.57
Inflation term premium	NaN	NaN	0.26	0.03	0.15	0.34

Table 8: Data and model moments. Columns 2 and 3 give the mean and standard deviation from U.S. Data over the sample 1985q1-2015q4. Columns 4 and 5 give the mean and standard deviation using simulated data from the model. Columns 6 and 7 give the mean by subsamples defined by the Taylor rate. Log output gap is computed as the log deviation from potential.The trend is constructed based on the HP filter using actual data.

	Subsample	
	$i^{TR} \leq -2\%$	$i^{TR} > 4\%$
Model	6.8	-2.4
Data: IS	9.6	2.0
Data: BE	12.6	2.5

Table 9: Slope coefficient in the regression of stock return on the change in breakeven in the model (top row) and in the data (using inflation swaps or breakevens).

(-4.5% average inflation in this subsample), larger than anything observed in the recent U.S. data. Hence, the model generates too much deflation, compared to what has been observed in modern economies. This is a well known limitation of the New Keynesian ZLB model. To put it another way, the model also implies that real interest rates fall significantly less than nominal interest rates, which is at odds with the data.

5.5 Comparative analysis

We first discuss the role of risk aversion, then the role of monetary policy, and finally the role of other parameters.

5.5.1 Role of risk aversion

There are two reasons why it is interesting to study how risk aversion affects our findings. First, our model is a standard New Keynesian model with the ZLB, but with high risk aversion. How do these "nonstandard" preferences affect the responses of consumption and inflation, which have been studied extensively in the New Keynesian literature in models with low risk aversion? Second, a broader debate exists over the importance of risk aversion for macroeconomic dynamics. Tallarini (2000) provides a quantitative example where risk aversion has little effect on business cycle dynamics. Risk aversion affects the steady-state level of the economy but responses to small shocks do not generate a change in risk, and hence risk aversion does not affect business cycle dynamics.

In our benchmark calibration, there is a small but significant effect of risk aversion on macro dynamics, but only when the economy is near the zero lower bound. The logic is as follows. When the economy hits the ZLB, macro volatility rises, chiefly because the effect of preference shocks on consumption and inflation becomes larger. This higher volatility in turn leads to higher precautionary savings which reinforce the recession. This effect is stronger with high risk aversion. As a result, we observe that inflation and consumption fall more when the economy becomes closer to the ZLB in the case of high risk aversion, than in the case of low risk aversion. Figure 7 depicts this results by comparing the impulse response functions in our benchmark calibration and in the case of $\alpha = 0$ (and we keep the other parameters at their benchmark values). The left panel shows that, as in Tallarini (2000), business cycle dynamics are completely unaffected if the economy is far from the ZLB (the two lines are on top of each other). The right panel shows that, at the ZLB, there is a meaningfully deeper recession in a model with high risk aversion.

5.5.2 Role of Monetary Policy

In our model,we assume that monetary policy follows a fairly standard Taylor rule. This is an empirically realistic depiction of monetary policy, and hence a reasonable benchmark for positive purposes. From a normative point of view, however, one can use the model to evaluate alternative monetary policy rules, which might stabilize output and inflation better. Absent the ZLB, our model features the so-called "divine coincidence", i.e. perfect inflation stabilization is possible, and optimal.[32] With the ZLB, optimal policy ought to respond aggressively, to

[32] This result continues to hold in our model despite the non-standard preferences. These preferences further imply that the welfare benefits of business cycle stabilization is significantly higher than in the standard New Keynesian model.

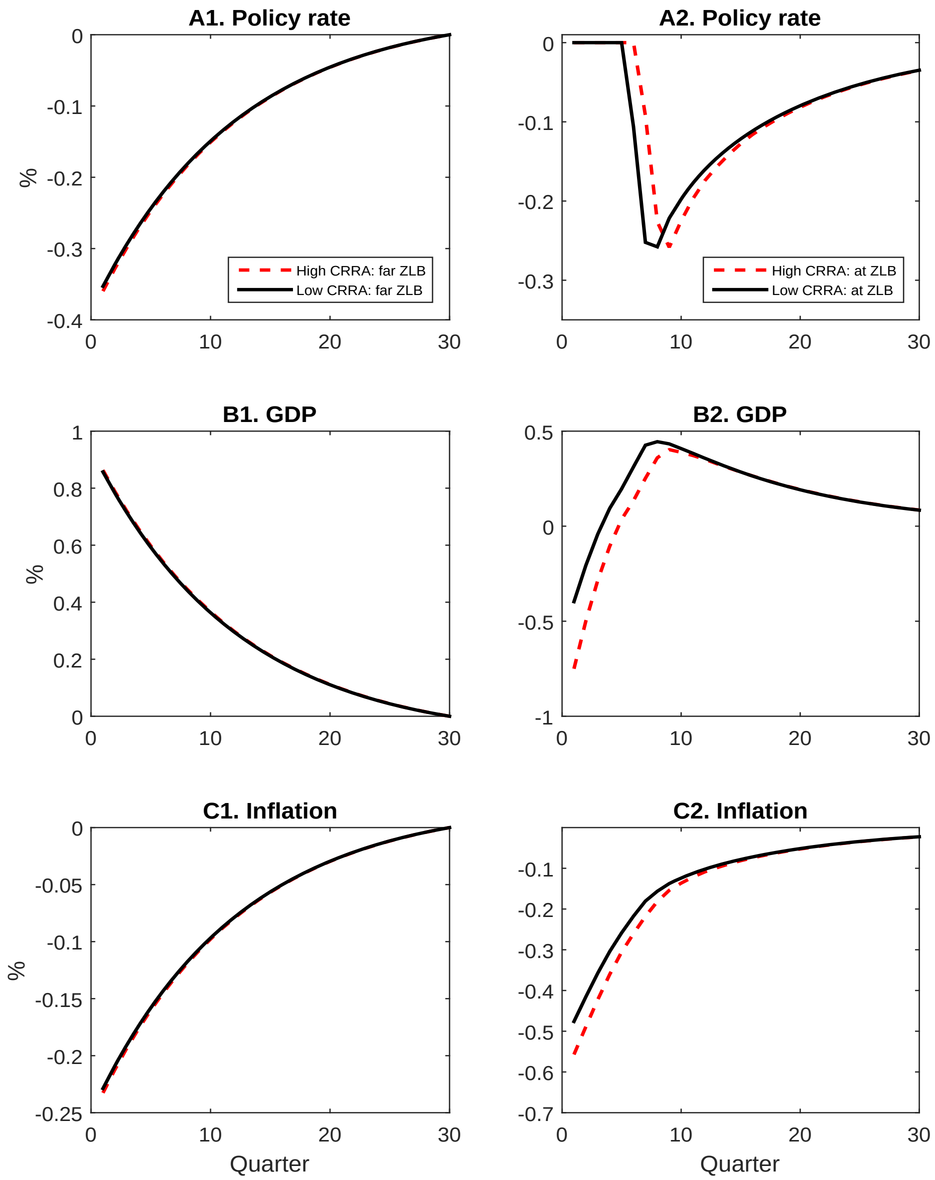

Figure 7: Impulse response function to a one-standard deviation productivity shock in the benchmark model (high risk aversion) and in the same model but with low risk-aversion, when the economy is at the ZLB vs. in steady-state (far from the ZLB). The solid black lines present impulse responses for low risk aversion, while the dashed red lines show the results for high risk aversion. The left panels are for the case at steady state (far from the ZLB). The right panels present impulse responses at the ZLB, which is calculated as the difference between two parths: (i) a path with only large liquidity shock that brings the economy to the ZLB, and (ii) a path with the same shock, plus a one-standard-deviation shock.

anticipate the possibility that the ZLB might bind (Adam and Billi (2006)).

Table 10 presents the results for different monetary policy rules. In our first experiment, we set the coefficient on output to 0. Note that inflation and output (rather than output gap) stabilization are equivalent for demand shocks, but not for supply shocks. As a result, when the central bank puts less weigh on stabilizing output ($\phi_y = 0$), it will stabilize inflation more in response to supply shocks. This leads to lower volatility of inflation (1.29% vs. 1.5%). Moreover, by focusing on inflation, the central bank is better able to stabilize the real output gap (1.29% vs. 1.5%). Both effects contribute to lower term premia.

Second, the weight of stabilizing inflation gap is raised substantially so that inflation gap can be stabilized almost completely. When we put more weight on stabilizing inflation ($\phi_\pi = 20$), inflation risk becomes very small (volatility of 0.10%) and hence inflation premia essentially disappear (0bps). The model never hits the ZLB owing to the strong response of the central bank. However, such a rule is probably unrealistic in practice because it amounts to neglecting output stabilization (e.g. in the face of a cost-push shock).

Third, we reduce the intercept of the Taylor rule, which is mechanically equivalent to a higher inflation target Π^*. As a result, we obtain higher average inflation (2.16% vs. 2.01%). There is little effect on term premia, however.

Finally, instead of using the Taylor rule that stabilizes GDP around its trend, we use a more conventional rule that stabilizes GDP around potential GDP. (Note that potential GDP is affected by technology level as explained in section 4.4.) The central bank stabilizes both output gap and inflation more efficiently in response to supply shocks, leading to lower volatility (1.2% vs. 1.5% and 1.15% vs. 1.5%). As a result, the level of term premia is lower.

6 Conclusion

We have shown that financial markets data suggest that inflation, which was historically associated with bad economic outcomes, became associated with good outcomes after 2008. A simple New Keynesian model that incorporates the zero lower bound can rationalize this finding qualitatively, and to some extent quantitatively. This changing correlation implies a reduction in the inflation premium as the economy becomes closer to the zero lower bound. We have also highlighted some significant limitations of this framework - the behavior of inflation, real interest rates, and term premia is at odds with the data. In that sense, this simple New Keynesian model is somewhat lacking, which motivates future research. One potential avenue is to extend the simple framework we used, by incorporating additional realistic frictions or a wider array

A. All observations

	Data		Benchmark		$\pi_y = 0$		$\phi_\pi = 20$		$R^* = 1.049$		Standard TR	
	Mean	Std	Mean	Std	Mean	Std	Mean	Std	Mean	Std	Mean	Std
logOutputGap	0.00	1.06	0.03	1.50	0.07	1.29	-0.02	0.10	0.02	1.58	0.03	1.20
logLaborGap	0.00	1.14	0.01	0.75	0.02	0.66	-0.01	0.48	0.00	0.77	0.01	0.65
π	1.80	1.20	2.01	1.50	1.93	1.29	1.98	0.10	2.16	1.48	1.99	1.15
$y^{n(1)}$	3.60	2.60	4.98	2.57	4.87	2.43	4.89	1.58	5.13	2.56	4.93	2.33
$y^{n(40)}$	5.50	2.30	5.35	0.77	5.13	0.70	5.10	0.44	5.51	0.77	5.19	0.67
$y^{r(1)}$	NaN	NaN	2.94	1.28	2.93	1.28	2.91	1.49	2.94	1.28	2.93	1.30
$y^{r(40)}$	3.30	0.90	3.12	0.32	3.11	0.32	3.11	0.37	3.13	0.32	3.11	0.33
$BE^{(40)}$	2.20	0.40	2.23	0.46	2.02	0.38	1.99	0.07	2.39	0.46	2.08	0.34
Real TP	NaN	NaN	0.18	0.02	0.18	0.01	0.22	0.01	0.19	0.02	0.19	0.01
Nominal TP	NaN	NaN	0.44	0.05	0.33	0.03	0.23	0.01	0.45	0.05	0.32	0.02
Inflation TP	NaN	NaN	0.26	0.03	0.14	0.02	0.01	0.00	0.26	0.03	0.13	0.01

B. Conditional on $i^{TR} < -2\%$

	Data		Benchmark		$\pi_y = 0$		$\phi_\pi = 20$		$R^* = 1.049$		Standard TR	
	Mean	Std	Mean	Std	Mean	Std	Mean	Std	Mean	Std	Mean	Std
logOutputGap	0.00	1.06	-2.25	0.84	-2.84	0.65	NaN	NaN	-2.73	0.67	-3.08	0.53
logLaborGap	0.00	1.14	-1.61	0.27	-1.69	0.34	NaN	NaN	-1.76	0.27	-1.87	0.25
π	1.80	1.20	-4.53	0.96	-4.41	0.92	NaN	NaN	-4.40	0.88	-2.55	1.07
$y^{n(1)}$	3.60	2.60	0.00	0.00	0.00	0.00	NaN	NaN	-0.00	0.00	0.00	0.06
$y^{n(40)}$	5.50	2.30	3.39	0.09	3.26	0.08	NaN	NaN	3.48	0.08	3.47	0.13
$y^{r(1)}$	NaN	NaN	3.73	0.84	3.58	0.80	NaN	NaN	3.58	0.77	1.96	0.89
$y^{r(40)}$	3.30	0.90	2.66	0.07	2.59	0.07	NaN	NaN	2.61	0.06	2.43	0.04
$BE^{(40)}$	2.20	0.40	0.72	0.16	0.68	0.14	NaN	NaN	0.87	0.14	1.04	0.16
Real TP	NaN	NaN	0.12	0.00	0.14	0.00	NaN	NaN	0.13	0.00	0.15	0.00
Nominal TP	NaN	NaN	0.28	0.01	0.21	0.01	NaN	NaN	0.29	0.01	0.23	0.01
Inflation TP	NaN	NaN	0.15	0.01	0.07	0.01	NaN	NaN	0.16	0.01	0.09	0.01

Table 10: Data and model moments for different parameters. In each case, we vary one parameter and keep the other parameters at the benchmark values. Columns 2 and 3 give the mean and standard deviation from the U.S. Data and columns 4 and 5 give the mean and standard deviation in the benchmark model. Columns 6-13 give the mean and standard deviation for different monetary policy rules: no response to output, a high response to inflation, a lower Taylor rule intercept, and a standard Taylor rule.The top panel reports the full sample and the bottom panel the subsample when the Taylor rate is less than -2%.

of shocks. Another avenue is to consider alternative economic mechanisms linking inflation and output or stock returns, such as the debt-deflation channel. Yet another possibility to generate a change in covariance from a shift in the structural parameters - e.g., in the shock process.

We also need to mention another limitation of our framework. In our model, the inflation premium has no direct effect on the economy. In reality, a lower inflation premium might stimulate the economy by lowering the interest rates faced by borrowers. It would be interesting to incorporate such feedbacks in the model.

Most central banks in developed economies now expect the zero lower bound to be a recurring constraint on policymaking. This means the effects we document and the channels we study will likely be at play again in the future.

References

Abel, A., 1999. Risk premia and term premia in general equilibrium. Journal of Monetary Economics 43(1), 3–33.

Adam, K., Billi, M., 2006. Optimal monetary policy under commitment with a zero bound on nominal interest rates. Journal of Money, Credit, and Banking 38(3), 1877–1905.

Adrian, T., Crump, R. K., Moench, E., 2013. Pricing the term structure with linear regressions. Journal of Financial Economics 110(1), 110–138.

Ai, H., Bansal, R., 2018. Risk preferences and the macro announcement premium. Econometrica 86(4), 1383–1430.

Ajello, A., Benzoni, L., Chyruk, O., 2014. Core and 'crust': Consumer prices and the term structure of interest rates. Working Paper, Federal Reserve Bank of Chicago.

Altavilla, C., Giannone, D., Modugno, M., 2017. Low frequency effects of macroeconomic news on government bond yields. Journal of Monetary Economics 92(C), 31 – 46.

Andrews, D. W., 1993. Test for parameter instability and structural change with unknown change point. Econometrica 61 (4), 821 – 856.

Ang, A., Bekaert, G., Wei, M., 2008. The term structure of real rates and expected inflation. Journal of Finance 63 (2), 797–849.

Anzoategui, D., Comin, D., Gertler, M., Martinez, J., 2017. Endogenous technology adoption and r&d as sourses of business cycle persistence. NBER Working Paper.

Arouba, B., Cuba-Borda, P., Schorfheide, F., 2018. Macroeconomic dynamics near the zlb: A tale of two countries. The Review of Economic Studies 85 (1), 87–118.

Bai, J., Perron, P., 1998. Estimating and testing linear models with multiple structural changes. Econometrica 66 (1), 47 – 78.

Bai, J., Perron, P., 2003. Computation and analysis of multiple structural change models. Journal of Applied Econometrics 18, 1 – 22.

Bansal, R., Shaliastovich, I., 2013. A long-run risks explanation of predictability puzzles in bond and currency markets. The Review of Financial Studies 26(1), 1–33.

Bernanke, B., Kuttner, K., 2005. What explains the stock market's reaction to federal reserve policy. Journal of Finance 60(3), 1221–1257.

Bianchi, F., Melosi, L., 2017. Escaping the great recession. American Economic Review 107 (4), 1030–58.

Bilal, M., 2017. Zeroing in: Asset pricing near the zero lower bound. Available at SSRN 2897926.

Boudoukh, J., Richardson, M., 1993. Stock returns and inflation: A long-horizon perspective. American Economic Review 83, 1346–1355.

Boyd, J., Hu, J., Jagannathan, R., 2005. The stock market's reaction to unemployment news: Why bad news is usually good for stocks. Journal of Finance.

Branger, N., Schlag, C., Shaliastovich, I., Song, D., 2016. Macroeconomic bond risks at the zero lower bound. Available at SSRN 2820207.

Breach, T., D'Amico, S., Orphanides, A., 2016. The term structure and inflation uncertainty. FRB of Chicago Working Paper No. WP-2016-22.

Campbell, J., Pflueger, C., Viceira, L. M., 2014. Monetary policy drivers of bond and equity risks",. NBER Working Paper.

Campbell, J., Vuolteenaho, T., 2004. Inflation illusion and stock prices. The American Economic Review 94 (2), 19–23.

Campbell, J. Y., Sunderam, A., Viceira, L. M., 2009. Inflation bets or deflation hedges? the changing risks of nominal bonds. National Bureau of Economic Research Working Paper.

Center, O. S., 1987. Ohio supercomputer center. `http://osc.edu/ark:/19495/f5s1ph73`.

Christiano, L., Eichenbaum, M., Rebelo, S., 2010. When is the government spending multiplier is large? Northwestern University Working Paper.

DAmico, S., Kim, D., Wei, M., 2018. Tips from tips: the information content of treasury inflation-protected security prices. Journal of Finance 53 (1), 395–436.

Datta, D., Johannsen, B. K., Kwon, H., Vigfusson, R. J., 2018. Oil, equities, and the zero lower bound. FEDS Working paper.

David, A., Veronesi, P., 2013. What ties return volatilities to price valuations and fundamentals?",. Journal of Political Economy 121 (4), 682–746.

Debortoli, D., Gali, J., Gambetti, L., 2019. On the empirical (ir)relevance of the zero lower bound constraint. Barcelona GSE Working Paper No. 1013.

Del Negro, M., Giannoni, M. P., Schorfheide, F., 2015. Inflation in the great recession and new keynesian models. American Economic Journal: Macroeconomics 7(1), 168–196.

Duarte, F., 2013. Inflation risk and the cross section of stock returns. Federal Reserve Bank of New York Staff Report.

Eggertsson, G., Krugman, P., 2012. Debt, deleveraging and the liquidity trap: A fisher - minsky - koo approach. Quarterly Journal of Economics 127(3), 1469–1513.

Eggertsson, G., Woodford, M., 2003. The zero bound on interest rates and optimal monetary policy. NBER Working Paper.

Epstein, L. G., Zin, S. E., 1989. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. Econometrica 57(4), 937–969.

Fama, E., Schwert, W., 1977. Asset returns and inflation. Journal of Financial Economics 5(2), 115–146.

Faust, J., Rogers, J. H., Wang, S.-Y. B., Wright, J. H., 2007. The high-frequency response of exchange rates and interest rates to macroeconomic announcements. Journal of Monetary Economics 54 (4), 1051–1068.

Fernandez-Villaverde, J., Gordon, G., Guerron-Quintana, P., Rubio-Ramirez, J., 2015. Nonlinear adventures at the zero lower bound. Journal of Economic Dynamics and Control 57, 182 –204.

Fischer, J. D., 2014. On the structural interpretation of the smets-wouters risk premium shock. Working Paper, Federal Reserve Bank of Chicago.

Fleckenstein, M., Longstaff, F., Lustig, H., 2014. Why does the treasury issue tips? the tips-treasury bond puzzle. Journal of Finance 69, 2151–2197.

Fleckenstein, M., Longstaff, F., Lustig, H., 2017. Deflation risk. The Review of Financial Studies 30, 2719–2760.

Gali, J., 2008. Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework. Princeton University Press.

Gilchrist, S., Schoenle, R., Sim, J., Zakrajsek, E., 2017. Inflation dynamics during the financial crisis. American Economic Review 107 (3), 785–823.

Gorodnichenko, Y., Weber, M., 2016. Are sticky prices costly? evidence from the stock market. American Economic Review 106(1), 165–199.

Gust, C., Herbst, E. Lopez-Salido, D., Smith, M., 2017. The emperical implications of the interest-rate lower bound. American Economic Review 107(7), 1971–2006.

Hordahl, P., Tristani, O., 2014. Inflation risk premia in the euro area and the united states. International Journal of Central Banking 2014(3), 1–47.

Kim, D. H., Wright, J. H., 2005. An arbitrage-free three-factor term structure model and the recent behavior of long-term yields and distant-horizon forward rates. Finance and Economics Discussion Series, Working Paper.

Krishnamurthy, A., Vissing-Jorgensen, A., 2012. The aggregate demand for treasury debt. Journal of Political Economy 120(2), 233–267.

Krugman, P., 1998. It's baaack: Japan's slump and the return of the liquidity trap. Brookings Papers on Economic Activity 2, 137–205.

Law, T. H., Song, D., Yaron, A., 2018. Fearing the fed: How wall street reads main street. Available at SSRN 3092629.

Leeper, E. M., Traum, N., Walker, T. B., 2017. Clearing up the fiscal multiplier morass. American Economic Review 107(8), 2409 – 2454.

Li, E., Palomino, F., 2014. Nominal rigidities, asset returns, and monetary policy. Journal of Monetary Economics 66, 210–225.

Marshall, D., 1992. Inflation and asset returns in a monetary economy. Journal of Finance 47(4), 1315–1342.

Miao, J., Ngo, P. V., 2014. Does calvo meet rotemberg at the zero lower bound?". Working Paper, Boston Univerisity and Cleveland State University.

Modigliani, F., Cohn, R., 1979. Inflation, rational valuation, and the market. Financial Analysts Journal 35(3), 24–44.

Nakata, T., 2017. Uncertainty at the zero lower bound. American Economic Journal: Macroeconomics 9 (3), 186–221.

Nakata, T., Tanaka, H., 2016. Equilibrium yield curves and the zero lower bound. Working Paper, Federal Reserve Board.

Ngo, V. P., 2014. Optimal discretionary monetary policy in a micro-founded model with a zero lower bound on nominal interest rate. Journal of Economic Dynamics and Control 45, 44–65.

Ngo, V. P., 2018. The risk of hitting the zero lower bound and the optimal inflation target. Macroeconomic Dynamics (Published online by Cambridge University Press: 02 May 2017) 22 (2), 402–425.

Palomino, F., 2012. Bond risk premiums and optimal monetary policy. Review of Economic Dynamics 15(1), 19–40.

Piazzesi, M., Schneider, M., 2006. Equilibrium yield curves. NBER Working Paper.

Rigobon, R., Sack, B., 2004. The impact of monetary policy on asset prices. Journal of Monetary Economics 51(8), 1553–1575.

Rotemberg, J., 1982. Sticky prices in the united states. Journal of Political Economy 90, 1187–211.

Roussellet, G., 2018. The term structure of macroeconomic risks at the zero lower bound. Available at SSRN 2863271.

Rudebusch, G., Swanson, E., 2008. Examining the bond premium puzzle with a dsge model. Journal of Monetary Economics 55, S111–S126.

Rudebusch, G., Swanson, E., 2012. The bond premium in a dsge model with long-run real and nominal risks. American Economic Journal: Macroeconomics 4(1), 105–143.

Song, D., 2017. Bond market exposures to macroeconomic and monetary policy risks. The Review of Financial Studies 30 (8), 2761–2817.

Swanson, E., 2015a. A macroeconomic model of equities and real, nominal, and defaultable debt. Working Paper, UC Irvine.

Swanson, E., 2015b. Risk aversion, risk premia, and the labor margin with generalized recursive preferences. Working Paper, Federal Reserve Bank of San Francisco.

Tallarini, T., 2000. Risk-sensitive real business cycles. Journal of Monetary Economics 45 (3), 507–532.

Taylor, B., 1993. Discretion versus policy rules in practice. Carnegie-Rochester Conference Series on Public Policy 39, 195–214.

Weber, M., 2015. Nominal rigidities and asset pricing. Working Paper, University of Chicago.

Wieland, J., 2014. Are negative supply shocks expansionary at the zero lower bound. Mimeo, UCSD.

Woodford, M., 2003. Interest and Prices: Foundations of a Theory of Monetary Policy. Princeton University Press.

Woodford, M., 2011. Simple analytics of the government expenditure multiplier. American Economic Journal: Macroeconomics 3, 1–35.

Wright, J. H., 2011. Term premia and inflation uncertainty: Empirical evidence from an international panel dataset. American Economic Review 101 (4), 1514–34.