BANK OF ENGLAND

Staff Working Paper No. 931

A tail of three occasionally-binding constraints: a modelling approach to GDP-at-Risk

David Aikman, Kristina Bluwstein and Sudipto Karmakar

January 2023

This is an updated version of the Staff Working Paper originally published on 16 July 2021

BANK OF ENGLAND

Staff Working Paper No. 931

A tail of three occasionally-binding constraints: a modelling approach to GDP-at-Risk

David Aikman,[1] Kristina Bluwstein[2] and Sudipto Karmakar[3]

Abstract

We build a semi-structural New Keynesian model to study the drivers of macroeconomic tail risk ('GDP-at-Risk'). Our model features three key non-linearities: a lower bound on nominal interest rates; a credit crunch in bank loan supply when bank capital depletes; and deleveraging by borrowers when debt service burdens become excessive. These non-linearities can interact to amplify GDP-at-Risk: for example, when debt burdens rise sufficiently, this increases the risk of debt deleveraging but also that of a credit crunch and hitting the effective lower bound. We simulate a persistent inflation shock to analyse how these interactions might operate at this juncture.

Key words: Financial crises, bank capital, debt deleveraging, macroprudential policy, effective lower bound, GDP-at-Risk.

JEL classification: E52, G01, G28.

(1) King's Business School, King's College London. Email: david.aikman@kcl.ac.uk
(2) Bank of England. Email: kristina.bluwstein@bankofengland.co.uk
(3) Bank of England. Email: sudipto.karmakar@bankofengland.co.uk

The views expressed are those of the authors and should not be taken to represent those of the Bank of England or its policy committees. For helpful discussions and comments, we would like to thank Tobias Adrian, Stephen Cecchetti, Thibaut Duprey, Charles Goodhart, Richard Harrison, Alejandro Van der Ghote and Matthew Waldron and seminar and conference participants at the Bank of England, Bank for International Settlements, RCEA Money Macro & Finance Conference, and the Annual Workshop of ESCB Research Cluster 3. Part of this paper was written while David Aikman was at the Bank of England.

The Bank's working paper series can be found at www.bankofengland.co.uk/working-paper/staff-working-papers

Bank of England, Threadneedle Street, London, EC2R 8AH
Email enquiries@bankofengland.co.uk

ISSN 1749-9135 (on-line)

1. Introduction

In recent years, the concept of 'GDP-at-Risk' has gained prominence as an aggregate macro measure of financial stability, as it captures the likelihood of extreme events in the tail of the distribution of real output like financial market stresses and crises. This approach opens up the possibility of aggregating different sources of risks into one summary metric which can help to add discipline to central banks' financial stability monitoring efforts.[1] If developments in the financial system can be quantified in the common currency of what they imply for tail risks to the economy, risks can be compared directly and mitigating efforts can target areas of greatest threat. There is also the prospect that such an approach might aid central banks' efforts to communicate risks to the public, and provide a quantitative anchor for calibrating macroprudential tools such as the countercyclical capital buffer (CCyB).

Previous studies have mostly focused on the empirical estimation of GDP-at-risk and analysed the drivers of tail risks via the contribution of individual variables or financial indices (e.g. Adrian et al., 2018; Aikman et al., 2019). While these empirical papers have found a strong role for financial factors in skewing tail risks, the economic mechanisms behind what causes the skewness in the GDP growth distribution have been less explored.

In this paper, we present a novel semi-structural, non-linear model of GDP-at-Risk. The model incorporates three occasionally binding constraints: first, an effective lower bound (ELB) on policy rates which reduces the capacity of the central bank to cushion shocks; second, a capital constraint on banks that could cause a credit crunch when bank equity is sufficiently depleted; and third, deleveraging by borrowers when the cost of servicing debt becomes excessive. The mechanisms modelled via these constraints have taken centre stage in developments in the global economy over the past decade and have featured prominently in macroeconomic research since the crisis.[2] A key contribution of this paper is that we are able to capture the joint impact of these constraints and their interaction simultaneously, in a framework that is both simple and transparent. The model is 'semi-structural' because, for simplicity's sake, we do not provide a micro-founded description of the behaviour of optimising agents; instead we focus on providing a plausible macro-level description of how certain nonlinear constraints interact with macro-financial dynamics to generate tail risk. However, most of the model is rooted in standard New Keynesian theory.

We present two key results. First, we show that the overall effect of these constraints is to generate a distribution of GDP that is significantly fat-tailed – a feature that is absent in standard linear New Keynesian models. Calibrating our model to the pre-COVID period and applying the Shapley value approach, we find that the effective lower bound has the strongest effect on tail risks,

[1]See Cecchetti and Schoenholtz (2018) for discussion of the concept and Adrian et al. (2018), Adrian et al. (2019), and Aikman et al. (2019) for recent empirical analysis. See Cecchetti (2006) and De Nicolo and Lucchetta (2012) for early expositions of this approach and Suarez (20121) and Cecchetti and Suarez (2021) for optimal policy design using GDP at risk.

[2]For example, on the effective lower bound, see Kiley and Roberts (2017) and Bernanke (2017); on credit crunch dynamics, see Gertler and Kiyotaki (2015), Zhu (2008), and Karmakar (2016); and on household leverage, see Mian and Sufi (2010).

while the capital requirement and the deleveraging constraints make only a modest contribution.[3]

Second, we find an interesting interplay between these constraints: an economy prone to debt deleveraging will experience substantially more frequent credit crunch and effective lower bound episodes; moreover, low capital buffers make ELB episodes more frequent and vice versa. In terms of policy implications, this implies that monetary and macroprudential policies can interact with each other. As a rule of thumb, we find each 50 basis point decline in monetary policy headroom requires additional capital buffers of 1% of assets or 2-2.5% points lower debt service ratios (DSR) in order to hold tail risk constant. However, that only holds up to a certain level of buffers and monetary policy headroom.

A key advantage of the modelling approach we propose is that it allows researchers to conduct 'what if' and counterfactual scenarios. We apply the model to explore two different scenarios: (a) a historical decomposition of GDP-at-Risk for the United Kingdom over the past two decades, and (b) a state of the world with persistent inflation. Historically, we show that the depletion in bank capital buffers and the build-up in private sector non-financial debt contributed to the worsening of the growth outlook in the run up to the Global Financial Crisis. In the post-crisis environment, the limited space for monetary policy to cushion adverse shocks exerts a material impact, leading GDP-at-Risk to remain elevated for several years after the crisis. Despite the elevated risk level, by 2014 the marginal contribution of the bank capital constraint to risk is approximately zero. This suggests the usable capital buffers by this date – 2.3% of assets or approximately 6.5% of risk-weighted assets (with an average risk weight of 35%) – were of an adequate magnitude to absorb credit losses that might occur during a tail event.

For inflation, the model suggests that higher inflation and interest rates increase financial stability risks in the near term, as higher rates put pressure on debt servicing costs leading to a greater risk of debt deleveraging by heavily indebted households. There is also a risk of higher loan defaults eroding banks' equity capital, which could lead banks to tighten lending conditions – although this effect is small in our model given the size of banks' capital buffers. In the medium term, financial stability risks recede as central banks have greater scope to cut interest rates in a stress, cushioning the effects of negative shocks.

The policy implications of these findings are three-fold. First, our results suggest that the macro-financial scenarios used in bank stress tests should be made more severe than 'normal', when interest rates are closer to the effective lower bound. Second, we find significant spillovers from the deleveraging constraint to the frequency of ELB and credit crunch episodes, suggesting there are wider macroeconomic stability benefits to reducing debt deleveraging risks. Third, our results have implications for the level of capital buffers required to maintain stability. We find the marginal benefits of increasing usable capital buffers beyond around 2.5% of assets is negligible in terms of reducing the 5% level of GDP-at-Risk over the next 3-to-5 years. There is an important caveat to this statement, however. If there are frictions to maintaining capital buffers at their current level –

[3]This result is dependent on the baseline calibration to reflect key features of the UK economy in the post-GFC and the pre-COVID period, such as low interest rates and high capital buffers. In Section 5, we present results from a scenario of high inflation and subsequent interest rate increases.

associated with chronic low profitability in the banking system, for example – then a longer-term perspective is required. In longer-run simulations, our model suggests that in a time of ultra-low interest rates, the marginal benefits of increasing capital buffers remain positive for buffers up to around 4% of assets or 11.5% of risk-weighted assets (with an average risk weight of 35%).

Related literature. A large literature has documented the decline of GDP volatility before the financial crisis of 2008 (see, e.g., Blanchard and Simon (2001); Giannone et al. (2008)). In contrast, the current emerging literature on GDP-at-Risk goes beyond the second moment of GDP growth and analyses the whole conditional distribution of GDP growth.

Existing research has, empirically, quantified how easy current financial conditions can lead to heightened tail risks in the medium term. Adrian et al. (2018) was one of the first papers to demonstrate the term structure of growth-at-risk. Using panel quantile regressions for 11 advanced and 10 emerging market economies, it showed that the conditional distribution of GDP growth depends on financial conditions, with growth-at-risk (GaR) – defined as growth at the lower 5th percentile. Aikman et al. (2019) also used a quantile regression approach to examine how the downside risk to growth over the medium term is affected by a set of macroprudential indicators. The authors find that credit booms, property price booms and wide current account deficits each pose material downside risks to growth at horizons of 3 to 5 years and that such risks can be mitigated by increasing the capitalisation of the banking system. Lloyd et al. (2021) extend this analysis to include global factors and find that foreign factors are a key driver of macroeconomic tail risks. Adrian et al. (2019) study the conditional distribution of GDP growth as a function of economic and financial conditions. They show that the lower quantiles of GDP growth vary significantly with financial conditions while the upper quantiles remain stable over time. They also argue that amplification mechanisms in the financial sector generate the observed growth vulnerability dynamics. Loria et al. (2019) study the entire forecast distribution of GDP growth-shifts in response to shocks and find that contractionary shocks disproportionately increase downside risk. Chavleishvi et al. (2021) construct a structural quantile VAR model to account for the direct and indirect interactions between financial vulnerabilities, financial stress and real GDP growth. However, using a quantile methodology on a fairly small dataset could have some drawbacks, e.g. Plagborg-Moller et al. (2020) critically evaluate the ability to predict more than the evolution of the conditional mean using quantile regressions.

While the empirical literature has made great progress, there has been far fewer theoretical contributions that could explain the drivers of tail risks from a conceptual point or allow for counterfactual or 'what-if' exercises. An exception to that is Adrian and Duarte (2017) which presents a parsimonious macroeconomic framework for incorporating financial vulnerabilities in monetary policy, although without modelling the causes of vulnerability explicitly. Starting from the standard New Keynesian models of Woodford (2003) and Gali (2008), vulnerabilities in the model give rise to a sharply negatively skewed unconditional output distribution, and therefore, captures movements in the conditional GDP distribution that correspond to the downside risk of growth. Our model adds to this emerging theoretical strand of research by explicitly modeling three key non-linearities that are analysed in DSGE models i.e., the debt deleveraging channel, the capital requirement channel and the effective lower bound.

The paper is structured as follows. Section 2 presents the model alongside the three occasionally-binding constraints. Section 3 presents the dataset employed and explains our calibration strategy together with an evaluation of the model performance relative to the data. Section 4 presents the results and analyse the importance of the individual constraints and their interaction. Section 5 uses the model to analyse an inflation scenario and conduct a historical decomposition of GDP-at-Risk for the United Kingdom. Section 6 concludes with directions for future research.

2. A simple model of GDP-at-Risk

We augment the standard New Keynesian model, similar to that described in Woodford (2003) and Galí (2008), to include nonlinearities associated with three occasionally-binding constraints: (a) an effective lower bound on interest rates, (b) the potential that banks may restrict lending sharply when their capital position becomes impaired, and (c) the potential that the households and non-financial corporates deleverage sharply when their debt service burdens become too large. We first outline the basic structure of the model, before explaining how these constraints are modelled.

The model consists of six equations, which govern the evolution of six endogenous variables: the output gap, y_t, the inflation rate, π_t, the nominal interest rate, r_t, the loan spread, s_t, the level of private nonfinancial sector credit, b_t, and the banking system's aggregate capital ratio, k_t. All variables are defined in terms of deviations from their respective steady states.

Aggregate demand, inflation and interest rates in period t are determined in conventional fashion by 'IS' and 'Phillips' curves and a Taylor rule respectively:

$$y_t = \theta^y y_{t-1} - \theta^r(r_t - \pi_{t-1} + s_t) + \epsilon_t^y + \epsilon_t^d \tag{1}$$

$$\pi_t = \beta^\pi \pi_{t-1} + \beta^y y_{t-1} + \beta^s s_{t-1} + \epsilon_t^\pi \tag{2}$$

$$r_t = \max\left[\bar{r}, (1-\phi^r)(\phi^\pi \pi_t + \phi^y y_t) + \phi^r r_{t-1} + \epsilon_t^r\right] \tag{3}$$

where ϵ_t^y, ϵ_t^π and ϵ_t^r are shocks to output, inflation and the interest rate respectively; ϵ^d is a deleveraging shock, discussed further in the next section. We assume these shocks are orthogonal and follow first-order autoregressive processes $\epsilon_t^i = \rho^i \epsilon_{t-1}^i + u_t^i$, where $u_t^i \sim iid(0, \sigma_i^2)$ for $i = y, d, \pi, r$. $\bar{r}$ is the feasible lower bound on the central bank's monetary policy rate.

Following Curdia and Woodford (2010, 2016) among others, we introduce a role for fluctuations in credit spreads in driving macroeconomic variables. Higher credit spreads push down on aggregate demand via the IS equation by increasing the interest rates facing households and firms wishing to borrow for consumption and investment. Higher spreads also enter the Phillips curve and act as an endogenous 'cost-push' shock, reducing the economy's productive capacity. This channel could be given a number of structural interpretations. In models where financial frictions exist between households and lenders, such as Curdia and Woodford (2010), a cost-push effect such as this derives from the impact of higher loan spreads in distorting households' labour supply decisions. In models where firms face binding credit constraints, such as Gertler and Karadi (2011), higher loan spreads lead to capital shallowing, reducing labour productivity. They also increase the cost of financing firms' working capital needs, as in Carlstrom, Fuerst and Paustian (2010).

The financial side of our model consists of equations for credit demand, borrowing spreads and the evolution of banks' capital ratios:

$$b_t = y_t + \gamma^b b_{t-1} - \gamma^r (r_t - \pi_{t-1} + s_t) + \epsilon_t^b + \epsilon_t^d \tag{4}$$

$$s_t = f^s s_{t-1} + f^b b_t - \tilde{f}^k k_{t-1} + \epsilon_t^s \tag{5}$$

$$k_t = \delta^k k_{t-1} - \delta^r \Delta r_t + \delta^s (r_{t-1} + s_{t-1}) + \epsilon_t^k \tag{6}$$

where ϵ_t^b and ϵ_t^s are exogenous shocks to credit demand and the loan spread respectively and follow first-order autoregressive processes $\epsilon_t^j = \rho^j \epsilon_{t-1}^j + u_t^j$, where $u_t^j \sim iid(0, \sigma_j^2)$ for $j = b, s$.

Real credit demand, which we interpret as reflecting both credit provided by banks and non-banks, including via financial markets, is increasing in contemporaneous output (with unit elasticity) and decreasing in the real interest rate inclusive of the credit spread, with persistence γ^b. The supply of finance is assumed to be influenced by the level of bank capital: if capital ratios are eroded, the supply curve shifts leftwards, as governed by the regime-dependent parameter $\tilde{f}^k$, increasing the borrowing spread, s_t. Borrowing spreads are also assumed to be subject to exogenous shocks, ϵ^s. We interpret these shocks are reflecting shifts in investor risk appetite and in the risk bearing capacity of key non-bank intermediaries in financial markets.

Finally, the banking system's capital-to-asset ratio is a key state variable of the model. We assume banks' capital ratios are increasing with a lag in the *level* of nominal interest rates plus spread, but decreasing in the *change* in nominal interest rates. This captures the notion that the banking system benefits in the short-run from a reduction in interest rates, but a period of sustained low rates impairs net interest income.[4] Banks also suffer unexpected write-offs ϵ_t^k on their asset portfolios, the intensity of which varies procyclically according to $\epsilon_t^k = \nu^y y_{t-1} + \rho^k \epsilon_{t-1}^k + u_t^k$.

2.1. Three occasionally-binding constraints

The model features three distinct thresholds associated with the effective lower bound on nominal interest rates, $\bar{r}$, the level of the banking system's capital ratio, $\bar{k}$, and the debt-service ratio (DSR) of the private non-financial sector, $\bar{dsr}$, giving a total of seven potential regimes governing the economy's dynamics.

Effective lower bound (ELB) regime. We assume the central bank's instrument for monetary policy,

[4] A richer version of this mechanism would capture a 'reversal rate' effect, by which, beyond a certain point, reductions in interest rates would contract the supply of bank finance (see Brunnermeier and Koby (2019))

the short-term nominal interest rate, is bounded below by $\bar{r}$, limiting the extent to which it can be reduced to stimulate economic activity. This is motivated by the findings in Romer and Romer (2017), where the authors show that the degree of monetary policy space prior to financial distress—that is, whether the policy interest rate is above the zero lower bound —greatly affects the aftermath of crises. We do not consider the potential for quantitative easing or other extraordinary monetary policies to be implemented when short-term nominal interest rates hit this ELB.

Capital regime. We assume credit supply becomes highly sensitive to banks' capital positions when their effective capital buffer is exhausted and k falls below the threshold $\bar{k}$. That is:

$$\tilde{f^k} = \begin{cases} f^{\tilde{k}H} & \text{if } k_{t-1} \leq \bar{k} \\ f^{\tilde{k}L} & \text{otherwise} \end{cases} \tag{7}$$

where $f^{\tilde{k}H} \gg f^{\tilde{k}L}$. We interpret $\bar{k}$ as the effective deleveraging point for the bank, which in general will lie above its minimum regulatory capital ratio, reflecting both precautionary behaviour on the part of a bank and constraints placed on it by market participants (Karmakar (2016)).

This assumption is consistent with models that view banks as profit-maximising subject to an occasionally-binding constraint on their capital ratio. For instance, Van den Heuvel (2009) studies a bank's choice of new lending and dividends to maximise its profits subject a regulatory constraint on its minimum capital. A bank in this model will choose to maintain a buffer of excess capital over and above its regulatory minimum to lower the expected cost of future capital inadequacy. The model predicts that the lending decisions of an abundantly-capitalised bank will be close to those of an unconstrained bank, with small changes in capital having little impact on new lending. By contrast, when excess capital is low, the bank will choose to cut new lending aggressively in response to shocks to write offs, even when the constraint itself remains formally slack. Other models that feature a nonlinear relationship between credit supply and bank equity capital include He and Krishnamurthy (2019), Holden, Levine and Swarbrick (2019), von Peter (2004) and Valencia (2014).

Deleveraging regime. We assume the private sector reduces spending sharply to pay down debt when its debt service burden – interest payments on debt outstanding relative to income – exceeds the threshold $d\bar{s}r$. A similar mechanism for businesses has been documented in Kiyotaki and Moore (1997) and for households in Mian and Sufi (2009). In our model, deleveraging occurs when:

$$DSR_t \equiv r_t + s_t + b_t - y_t \geq d\bar{s}r \tag{8}$$

If borrowers attempt to delever by reducing spending en masse, and assuming, as is plausible, that borrowers have a higher marginal propensity to consume than savers, there will be an aggregate demand externality. We capture this by introducing a shock, ϵ_t^d, which reduces both credit demand and overall aggregate demand in period t by the same amount. The decline in income initially frustrates the attempt to reduce the debt burden. But over time, the stock of debt falls by more than income in response to this shock because the long-run income elasticity of credit demand exceeds one.

We assume the size of the deleveraging shock in period t is given by the gap between DSR_{t-1} and $\bar{dsr}$. That is, the private sector will aim to reduce its debt level by the amount required to return the debt service ratio to the threshold in period t, holding interest rates and overall output fixed:

$$\epsilon_t^d = \begin{cases} \bar{dsr} - (r_{t-1} + s_{t-1} + b_{t-1} - y_{t-1}) & \text{if } DSR_{t-1} \geq \bar{dsr} \\ 0 & \text{otherwise} \end{cases} \tag{9}$$

Moreover, we will treat deleveraging as a persistent process that continues to depress the economy for mutliple periods after it occurs: $\epsilon_t^d = \rho^d \epsilon_{t-1}^d$.

3. Calibration

3.1. Data set

For most of the analysis in the paper, we use six quarterly data series from 1992-2014: real GDP, inflation, Bank Rate, total credit (capturing borrowing by the non-financial private sector, from banks and non-banks), lending spreads, and the capital ratio of the UK banking system. Figure A.1 in the Appendix plots the raw time series used in the analysis.

Our measure of bank capital is the tangible common equity ratio (TCE), which is tangible common equity normalised by total tangible assets. Tangible common equity is defined as common equity minus preference shares and intangible assets.[5]

Our series on lending spreads is a weighted average of spreads on three different funding categories: household mortgage credit, consumer credit and corporate credit. The weights, which are consistent with those reported in the February 2014 Inflation Report of the Bank of England, are as follows: household credit is attributed a weight of 0.7 while corporate credit is assigned a weight of 0.3; within household credit, mortgage spreads have a weight of 0.75 (comprising deposit spreads 0.5, low-LTV spreads 0.21, and high LTV spreads 0.04) while consumer credit spreads have a weight of 0.25. Within corporate credit, spreads on small and medium enterprise loans have a weight of 0.37 while those on bond issues by large businesses are attributed a weight of 0.63.

For the purpose of calibrating model parameters, we demean all variables and, with the exception of the loan spread series, pass them through a one-sided Hodrick-Prescott filter with a smoothing parameter of 1,600.

3.2. Calibration approach

Table 1 records calibrated values of the model's parameters. We fix parameters of the Taylor rule and the slope of the Phillips curve to match posterior mean estimates used in the the Bank of England's main macro-forecasting model, COMPASS (as reported in Burgess et al. (2013)). To ensure the model

[5]The TCE measure is strongly correlated with other measures of banking system leverage. It has a correlation of 0.75 with the Bank of England's leverage indicator for the United Kingdom.

is dynamically stable, we set the endogenous persistence parameters to 0.5, with the exception of inflation persistence, which is set at 0.35.[6][7]

We choose $f^{\tilde{k}H}$, the parameter governing the strength of the bank credit crunch, by assuming that when the capital constraint binds the banking system attempts to keep its leverage ratio constant. This implies that each £1 loan loss is met with a $\frac{1}{\bar{k}}$ contraction in bank credit supply. That is, if $\bar{k} = 5\%$, the banking system will reduce its lending by 20% ($= 1/0.05$) for each % point reduction in its leverage ratio. In mapping this to $f^{\tilde{k}H}$ we further scale this parameter by 50% to reflect the fact that banks account for approximately half of total UK financial system assets (Bank of England, 2016).[8]

We chose other parameters to match certain target moments reported in the data and the literature. The exact matching exercise and target moments are reported in A.1.

The calibration for the shock process parameters are reported in Table A.2 in the Appendix. We first generate 5000 simulations from the model of 140 quarters, discarding the first 40 quarters from each. This gives us 5000 samples of simulated data of equivalent length to our dataset.

Each simulation is initialised at the model's steady state. We then compute a variety of statistics from each simulated path, and take the average across all 5000 samples. We calibrate the thresholds relating to the effective lower bound, the credit crunch the debt-deleveraging regimes as follows. We set $\bar{r} = -3.5$, consistent with an effective lower bound of 0% and an average value of Bank Rate over the sample of 3.5%. We set $\bar{k} = -1.5$, reflecting the idea that the UK economy experienced a credit crunch between 2007-2009, a period when our measure of the banking system's leverage ratio was 1.5% points below its mean. Following similar logic, we set $\bar{dsr} = 10$ as our measure of the debt-service ratio was 10% above its sample mean in these years.

We set the persistence parameter of each shock such that the model's autocorrelation function matches that of the data – see Figure A.3 in the Appendix presents the results of this exercise. We set the variance of the shock processes by attempting to match the standard deviations of each variable with those in the data; in doing so we also put weight on the 5th percentiles of model variables

[6]Consider the model in companion form $Y_t = FY_{t-1} + Gz_t$ where Y is the vector of n endogenous variables; $F \equiv A^{-1}B$ and $G \equiv A^{-1}C$ where A is the nxn (invertible) matrix of contemporaneous coefficients, B is the nxn matrix of lagged coefficients, and C is the matrix of loadings on the shock processes. The restrictions on F required for dynamic stability are that the moduli of the eigenvalues of F all lie within the unit circle.

[7]In general, the model can become unstable in these longer simulations. When monetary policy is constrained at the ELB, banks are restricting credit and debt deleveraging is underway, there is no stabilising force and the economy can be trapped in a situation where unboundedly bad outcomes are possible. To avert this, we assume a bank recapitalisation plan is implemented when the banking system's leverage ratio approaches 0% (falls 5% pts below its steady state level in our baseline calibration). This could be interpreted as a bail-in of the banking system's private creditors or as a taxpayer-funded equity injection. We calibrate this package by assuming it wipes out the current level of bad debt and pushes the system's leverage ratio above the credit crunch threshold. That is, if rc_t is the recapitalisation package and $\bar{\bar{k}}$ is the leverage ratio at which the plan is implemented, we have $rc_t = \bar{k} - \bar{\bar{k}} - \nu^y y_{t-1}$. Figure A.6 in the appendix illustrates this mechanism.

[8]The macroeconomic implications of this calibration can be compared with Guerrieri et al. (2015), who compare the impact of bank capital shocks across five dynamic general equilibrium models. The transfer shock they consider is worth the equivalent of 7.5% of GDP. Calibrating this for the UK (GDP c. £2tr; Core Tier 1 capital c. £350bn) means a shock of around £150bn, which would wipe out around 40% of banks' capital base. With an average leverage ratio of 3.8% in our dataset, this means a shock to ϵ^k of 1.5% points. Output in our model falls by 2.6% in response to this shock, with credit spreads increasing by 480 bps. The average peak responses reported in Guerrieri et al. (2015) are -2.4% for output and 130 basis points for credit spreads, with ranges in each case of (-0.5%,-5%) and (0,350).

Table 1: *Calibration of the model's parameters*

Parameter	Description	Value
Macro block:		
θ^y	Persistence of aggregate demand in the IS curve	0.5
θ^r	Interest elasticity of aggregate demand	-0.45
β^π	Persistence of inflation in Phillips curve	0.35
β^y	Slope of the (structural) Phillips curve	0.3
β^s	Impact of spreads on inflation	0.1
ϕ^π	Taylor rule coefficient on inflation	1.497
ϕ^y	Taylor rule coefficient on output gap	0.1512
ϕ^r	Taylor rule coefficient on inflation	0.6
Financial block:		
γ^b	Persistence of credit demand	0.5
γ^r	Interest elasticity of credit demand	-0.1
f^s	Persistence in loan spread	0.5
f^b	(Inverse) elast of credit supply wrt loan spread	0.05
$\tilde{f}^k$	Elasticity of loan spread wrt bank capital ratio	0.07 (baseline, f^{kL}) $0.5 f_b / k_{t-1}$ (credit crunch, f^{kH})[a]
δ^k	Persistence of bank capital ratio	0.5
δ^r	(Semi) elast of capital wrt change in nominal interest rate	0.05
δ^s	(Semi) elast of capital wrt level of nominal interest rate	0.01
ν^y	Elasticity of banks' write-offs wrt to output	0.033

Notes: This table presents calibrated values of the model's parameters. (a) We bound $f_{kH} \leq 0.5 f_b / 0.02$ to avoid bank deleveraging becoming unboundedly large as k approaches 0.

versus those in the data, plus correlations of endogenous variables with output in the model and data.

We centre the results around a baseline simulation in which the economy's steady state values of nominal interest rates, bank capital buffers and private nonfinancial sector debt imply the following distances to respective constraints:

– *Effective lower bound:* The economy has 3% points of headroom above the effective lower bound, consistent with an estimate of R^* of around 1% (Bank of England (2018)), a 2% inflation target and a 0% floor;

– *Capital requirements:* There is 2% points of leverage ratio headroom before the banking system delevers sharply. Given that the UK banking system had a leverage ratio of 5.4% in 2017Q4 in our dataset, this implies banks will rapidly delever at a leverage ratio of 3.4%. With an average risk weight of 35%, our calibration implies the banking system can absorb losses of almost 6% of risk-weighted assets (i.e. $\frac{2}{0.35}$) before sharp deleveraging occurs;

– *Debt deleveraging:* There is capacity for a 10% increase in debt service burden before private sector debt-deleveraging occurs.

3.3. Model performance

Table 2 reports the comparison between the model moments with the data. Overall, the standard deviations of GDP, inflation, the loan spread and credit are all close to those in the data. However, we overestimate those of Bank Rate (1.69 vs 1.05) and bank capital (0.63 vs 0.28). The 5th percentiles of Bank Rate is 0.15 in the model versus 0.4 in the data; while that of bank capital is 2.5% in the model versus 2.0% in the data.

The table also shows correlations of these variables with output. The model matches the cyclicality of Bank Rate, loan spreads and credit well. However, the correlations it generates of inflation and bank capital with output are not close. Inflation is uncorrelated with our measure of detrended output empirically, whereas our model sees these variables as intrinsically linked via the Phillips curve. Consistent with empirical evidence reported in Adrian and Shin (2010), the banking system leverage ratio has a negative correlation with output in the data reflecting the fact that banks operated with ever decreasing capital cushions in the boom prior to the crisis and were forced by regulators to raise capital significantly during the slump. In the model, bank capital is instead procylical – credit losses in the model are declining in output, while net interest income is increasing in output reflecting the behaviour of interest rates.

3.4. Long-run stochastic simulations to explore determinants of GDP-at-Risk

We begin the description of the model's properties by analysing long-run stochastic simulations. This allows us to explore the dynamics of the model in a general setting where the impact of

Table 2: *Comparison of the model with UK dataset*

Description	Data	Model
Standard deviations of model variables:		
Std. dev. of output (%)	1.30	1.18
Std. dev. of inflation (% pts)	0.82	0.57
Std. dev. of Bank Rate (% pts)	1.05	1.69
Std. dev. of loan spread (% pts)	0.81	0.77
Std. dev. of credit (%)	3.18	3.53
Std. dev. of bank capital (% pts)	0.28	0.63
Correlations with output:		
Inflation	-0.04	0.37
Bank Rate	0.66	0.42
Loan spread	-0.12	-0.20
Credit	0.19	0.35
Bank capital	-0.35	0.36

Notes: This table compares the standard deviations and output correlations in the data versus those in our model simulations.

particular assumptions about initial conditions is diminished.[9]

To characterise tail risk, we focus on the q-th percentile of GDP, which we refer to interchangeably as $GaR^q(y)$ or the q-% level of GDP-at-Risk. To measure this, we simulate the model a large number of times and report the average value across simulations. That is, if N is the number of simulations, the 5% level of GDP-at-Risk would be given by:

$$GaR^5(y) \equiv \sum_{n=1}^{N} \frac{1}{N} q^5(y_n) \tag{10}$$

where $q^5(y_n)$ is the 5th percentile of output in simulation $n \in N$. We have also explored an alternative measure of tail risk – the 'expected loss' of output, defined as the expected value of output conditional on output being below its q-th percentile. The results we report are qualitatively unaffected by which metric we focus on.

Figure 1 plots the density of output in these simulations against a version of the model with all constraints switched off ('linear model'). The distribution of output in the model model is unimodal with a heavy left-hand tail: the probability of large declines in output is significantly greater than that in the symmetric Gaussian distribution generated by the linear model: the 5% level of GDP-at-Risk is -2.8% in the baseline model vs -1.7% in the linear model. It compares to a 5th percentile in the historical data used to calibrate the model of -2.9%. In this simulation, the ELB

[9]Figures A.4 and A.5 in the appendix present impulse responses from the baseline version of the model alongside the linear no-constraints benchmark, which will give readers a further information about how the model behaves.

binds 11.1% of the time, the bank capital constraint binds 1.8% of the time and the DSR constraint binds 1.9% of the time. The model is therefore able to replicate the skewness of GDP growth that is found in the data.

Figure 1: *The density of GDP*

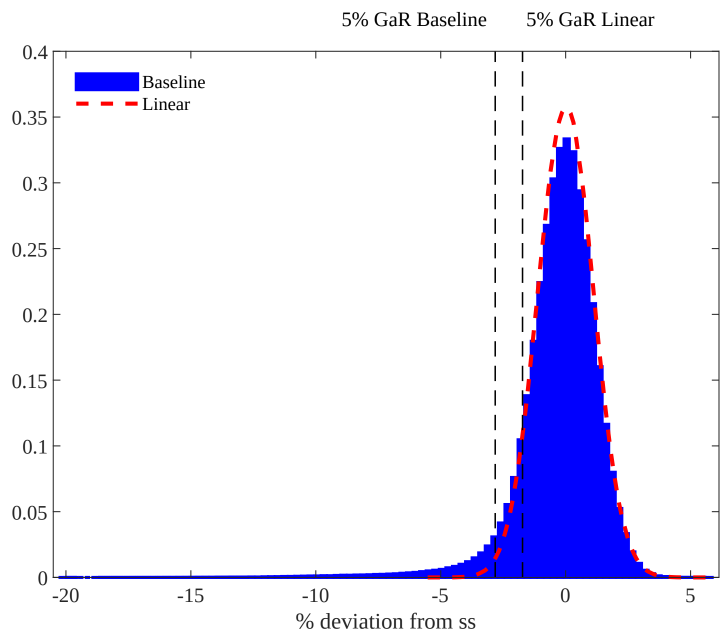

Notes: The figure presents the density of output in the baseline model versus that in a linear benchmark economy where are constraints have been switched off. It shows the kernel density estimate of the distribution of output across these 5000 one-hundred year simulations of the model.

4. The importance of non-linearities and their interaction

In this section, we delve further into the drivers of GDP-at-Risk by analysing the role of the three constraints. In particular, we ask three questions: First, what is the contribution of each to tail risk in our baseline calibration? Second, how do the constraints interact with one another? If one becomes more proximate, what impact does that have on the others? Third, how substitutable are the constraints from a risk perspective? We take each question in turn.

4.1. How important is each constraint?

Given the nonlinear structure of our model, in general the impact of any constraint will depend on the state of the economy, including the proximity of the other constraints considered. A simple approach of adding up the standalone impact of each constraint will therefore not adequately reflect their contributions to the overall risk level. Similarly, an approach of turning on constraints sequentially and recording their marginal impact will yield contributions that are sensitive to the order in which this is done.

Recognising this, we proceed by applying a methodology developed in game theory under which the contribution of a player to an overall outcome is given by her average marginal contribution across all possible subgroups in which she participates. In our context, implementing this 'Shapley value' approach (Shapley (1953)) amounts to simulating all combinations in which constraints are applied and assigning each constraint its average contribution across these combinations – see Appendix B for details.

Figure 2: *Contributions to GDP-at-Risk in the baseline model*

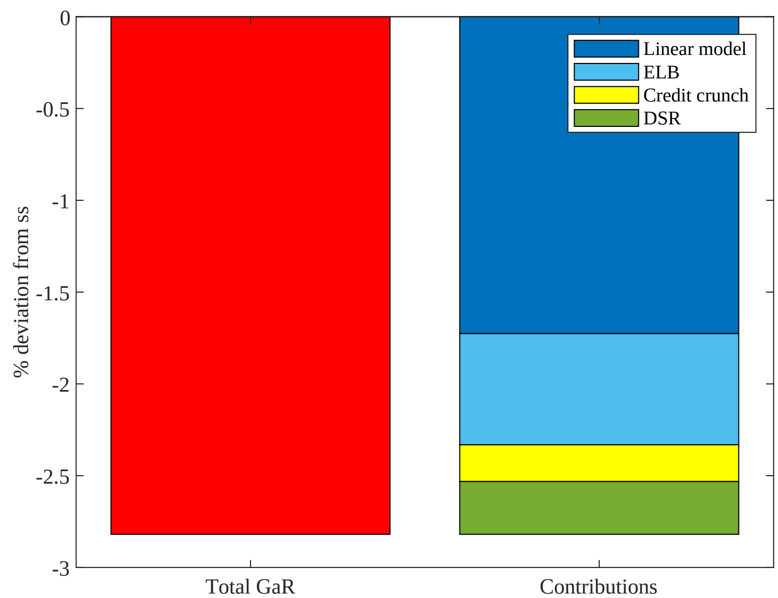

Notes: This figure presents Shapley value contributions of the three nonlinear constraints to the model's 5% level of GDP-at-Risk generated across 5000 one-hundred year simulations. See Appendix B for details of the method.

Figure 2 presents the results of this exercise. The left-hand column (red) is the 5% GDP-at-Risk level generated by our model over a large number of one-hundred year simulations – a risk level of -2.8%. The right-hand column shows the Shapley-value contributions of the three nonlinear constraints, plus the baseline linear model. The effective lower bound makes the largest contribution to tail risk, accounting for a full 0.6% points of the fall in output at the fifth percentile (shown in light blue). The contributions of the deleveraging constraint and bank capital constraint are 0.3% points (green) and 0.2% points (yellow) respectively. These contributions are all incremental to tail risk in the linear model of -1.7% (dark blue).

These contributions are sensitive to the assumed proximity of each constraint. In particular, the relatively smaller contribution of the bank constraint reflects the fact that bank capital buffers are currently large by historical standards. In section 5.1, we use our model to explore how these contributions have fluctuated in importance over recent UK macro-financial history.

4.2. Interactions between constraints

Given the structure of our model, when we change the proximity of one constraint, its impact on the overall risk level will in part reflect how it interacts with other constraints. For instance, a reduction in monetary policy headroom vis-a-vis the ELB will amplify the severity of credit crunch episodes and so part of the impact of the ELB constraint on the overall risk level will in part reflect this spillover.

Figure 3 provides a method for gauging the magnitude of these *cross-constraint spillovers*. The vertical axis on each plot shows the proportion of quarters (in %) with which each constraint binds on average across simulations. The horizontal axis shows the proximity of the bank capital, DSR and ELB constraints.

The top row shows the impact of changing headroom over the ELB constraint. As monetary policy headroom shrinks, the frequency of ELB episodes unsurprisingly increases. There is a material spillover to credit crunch episodes, as shown in the right-hand panel – intuitively when the capacity of the central bank to cushion falls in output diminishes, the result is deeper recessions and hence larger credit losses for banks. There is no impact on deleveraging risk, however. The reason for this is that in a credit crunch, the model predicts a rapid fall in the stock of outstanding credit, alleviating debt service pressure.

The second row performs the same exercise vis-a-vis the bank capital constraint. As buffers shrink, we observe an increase in the frequency of credit crunches. The relationship is highly nonlinear: the marginal impact of an erosion in capital buffers is greater when buffers are already low, with an inflexion point when usable capital falls to around 1.5% of assets (c. 4.5% of risk-weighted assets). The spillovers created are modest, however. The frequency of ELB episodes begins to increase once buffers hit 2% of assets. But there is no impact on the frequency of debt deleveraging episodes. The reason for this is that while loan spreads jump and income falls in credit crunch episodes, the impact on debt service is dominated by falls in Bank Rate and in the stock of credit.[10]

The bottom row examines spillovers created by the debt-deleveraging threshold. As the DSR constraint becomes more proximate, debt delevereraging frequency increases. Cross-constraint spillovers are substantial in this case – credit crunch and ELB episodes both increase in frequency significantly. The reason is straightforward: recessions become more frequent and more severe when the economy operates closer to the DSR constraint. So a given lower bound on interest rates will be hit more frequently, and banks will suffer larger credit losses, depleting their capital to the credit crunch threshold more often.

To summarise, low capital buffers make ELB episodes somewhat more frequent and low equilibrium real interest rates make credit crunch episodes somewhat more likely, but an economy more prone to debt deleveraging will experience substantially more frequent credit crunch and ELB episodes.

[10]A richer model would capture slow adjustment of the stock of credit in such circumstances, creating the potential for a larger spillover from capital buffers to debt deleveraging risk.

Figure 3: *Frequency of each constraint binding*

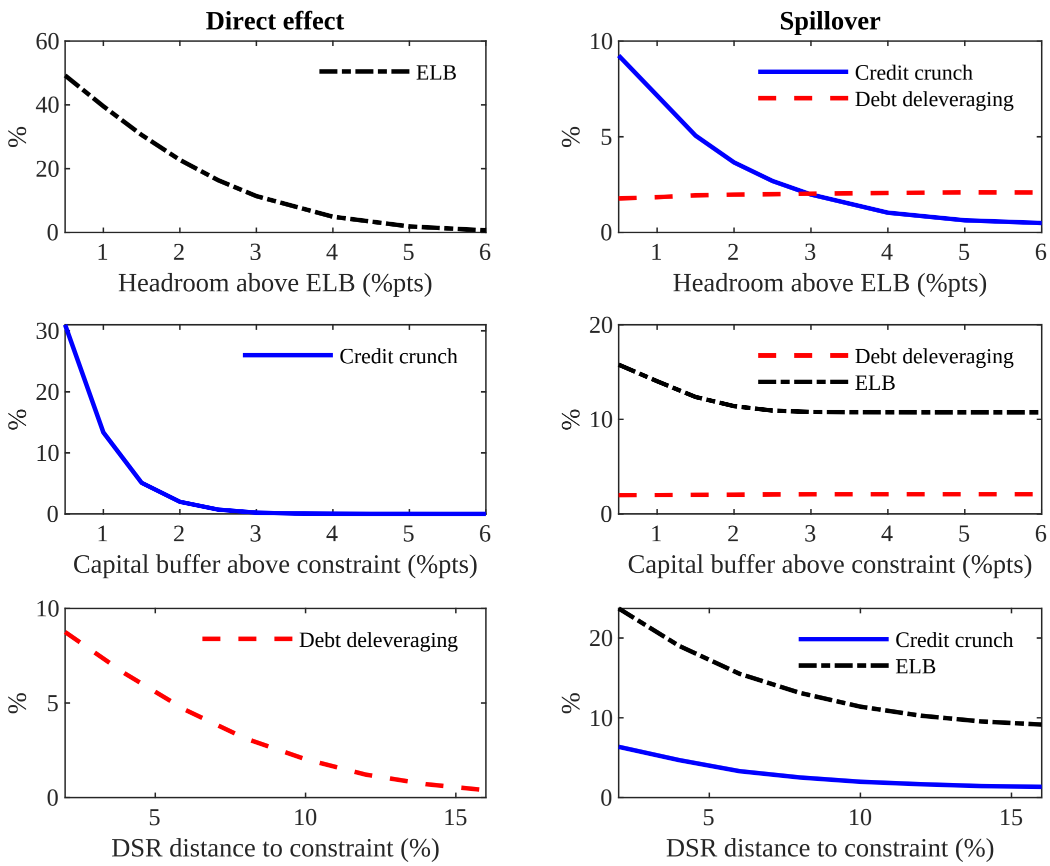

Notes: The figure plots the frequencies with which constraints bind as a function of the economy's headroom over its constraints. We simulate the model for 440 quarters, discarding the first 40 and repeating the process 1000 times. We report in the y-axis the average % of simulated quarters for which each constraint binds. The x-axis is the degree of headroom vis-a-vis the stated constraint assumed in the steady state. The first column ('Direct effect') shows the impact of making each constraint more proximate on the frequency with which it binds. The second column ('Spillover') shows the impact on the frequency of other constraints binding.

An immediate policy implication is that the returns to addressing vulnerabilities that lead to debt delevering (e.g. via borrower-based macroprudential measures) can be material – doing so reduces both debt deleveraging risk itself and the risk of pathologies associated with bank credit crunches and the ELB. Another implication relates to the design of bank stress tests. All else equal, these tests should be made more severe the greater is the proximity of the ELB and the larger the risk of debt-deleveraging.

4.3. Policy trade-offs and the 'marginal rate of risk substitution'

We conclude this section by assessing the extent to which the impact of a change in the proximity of one constraint can be offset by making another more distant – a concept we refer to as the *marginal rate of risk substitution*. Figure 4 illustrates this concept via 'iso-risk' curves – that is, combinations of constraints that generate the same level of GDP-at-Risk. We focus on the marginal rate of risk transformation between interest rate headroom and the other two constraints, holding one of which constant at its benchmark level in each plot.

The iso-risk curves in the upper panel slope downwards, reflecting the fact that diminished interest rate headroom must be accompanied by larger capital buffers if the overall risk level is to remain constant. Within the vicinity of our baseline calibration (interest rate headroom of 3% pts, capital buffers of 2% of assets), the slope of the curve is approximately -2, implying as a rule of thumb that a 1% pt increase in capital buffers is required to 'offset' the impact of interest rate headroom declining by 50bps. The contours become near-vertical once buffers exceed 4% of assets: capital buffers at this level are large enough to absorb losses so increases beyond this have no attenuating impact on risk. It is notable that the contours become more tightly spaced in the north-east region of the plot, indicating that the risk level is increasing more steeply in this region.

The lower panel illustrates the marginal rate of risk substitution with respect to the debt service constraint. Here, the lines slope upwards as, for a given risk level, the economy requires larger headroom vis-a-vis the DSR constraint when interest rate headroom falls. Their convexity indicates that DSR headroom becomes less effective as a risk-mitigant as that headroom increases. In the vicinity of the baseline calibration (DSR headroom of 10%, interest rate headroom of 3% points), the slope of the line is around 4-5, indicating that a 2-2.5% increase in DSR headroom is required to 'offset' the risk impact of interest rate headroom declining by 50bps.

5. Model applications

5.1. Analysing the drivers of GDP-at-Risk in recent decades

In this final section, we use the model to produce an estimate of fluctuations in GDP-at-Risk in the pre-and post-crisis period, alongside the role played by the constraints we consider. The exercise serves two purposes: first it allows us to gauge the relative importance of different factors in the lead-up to the global financial crisis; second, it provides a useful plausibility check on the behaviour of the model.

Figure 4: *Iso-risk curves*

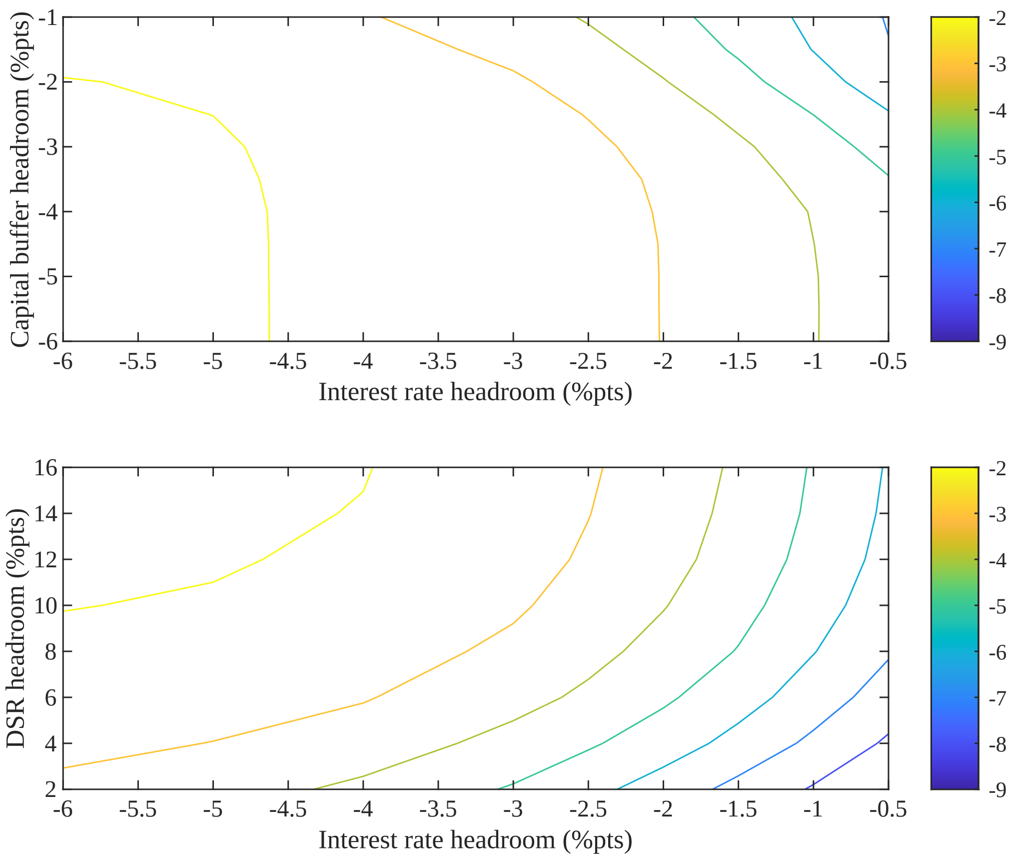

Notes: The figure plots combinations of constraint settings that deliver the same level of GDP-at-Risk. The upper panel sets out the marginal rate of risk transformation between interest rate headroom and the level of capital buffers for a given level of DSR headroom (assumed to be 10%, as per the benchmark calibration). The lower panel does likewise for interest rate headroom and headroom vis-a-vis the debt-service constraint for a given level of capital buffer (assumed to be 2% of assets, as per the benchmark calibration). The colourbar shows the corresponding level of GDP-at-Risk at which each contour is drawn, ranging from -2% to -10%.

To proceed, we first express each threshold in terms of the *level* of the relevant variable, allowing us to capture changes in the proximity of each constraint over the period. For the effective lower bound, we assume a threshold of 0% for the nominal interest rate. To calibrate the bank capital threshold, we appeal to the time series for banks' leverage ratios and choose a threshold consistent with the UK economy as having experienced a credit crunch during the global financial crisis. This implies setting the leverage ratio threshold at 2.3%. We take a similar approach for calibrating the debt-deleveraging threshold. We choose a DSR threshold value of 20%, consistent with debt deleveraging having begun in fourth quarter of 2007.

To map from model projections – which describe cyclical fluctuations – into level projections, we add in the trend component for each series. For example, to assess whether the bank capital constraint binds n periods ahead, we check whether $k^c_{t+n} + k^T_{t+n} \leq \bar{k}$, where the k^c_{t+n} is the output of the model. The trend component, which is estimated via a HP filter as described in section 3, is assumed to remain unchanged in the projection, $k^T_{t+n} = k^T_t$. For the debt service ratio, we introduce a measurement error term to explain the deviation between the model-implied DSR and the cyclical component of the DSR time series. We use an AR(1) model to project this measurement error forwards in our simulations, where the persistence parameter is estimated to be 0.92.

We initialise the model using realised data outturns in 1999Q4 and simulate 20 quarters forward, repeating this process 5000 times.[11] Our estimate of projected GDP-at-Risk for 1999Q4 is then given by the average 5th percentile of output across these simulations between quarters 12 and 20. We then iterate forward one quarter to 2000Q1 and repeat the process. We hold the distribution of shocks fixed at each iteration – this reflects the notion that shocks are unpredictable events, whereas changes in the state of the economy are observable and can be conditioned on.

Figure 5 presents the time series of UK GDP-at-Risk obtained by applying this method. The red line presents the full model-generated GDP-at-Risk projection at each date, alongside contributions of the effective lower bound, credit crunch and debt deleveraging constraints. This is the worst recession that might reasonably be expected over the period (up to a given confidence level); it is not a forecast of the most likely path of the economy. The linear model (dark blue bars) provides a useful benchmark for this exercise – it gives the *normal* level of tail risk we might expect if the economy were sufficiently far from the constraints that they can effectively be ignored. This is constant at -1.8% because projections from linear model at this horizon are unresponsive to initial conditions, the effects of which die away.

Viewed through the lens of the model, the risk outlook began to deteriorate in the early 2000s, driven by an erosion in banks' capital buffers and an increasing debt burden in the private nonfinancial sector. The likelihood of hitting the bank capital threshold and of breaching the debt service burden threshold therefore grew, as per the rising contributions of the credit crunch and debt deleveraging bars until the global financial crisis. By the first quarter of 2007, a time when the UK economy was booming, the model projects a 5% GDP-at-Risk level over the subsequent 3-5

[11] The DSR series we use is from the Bank of International Settlements, and it covers debt service payments of the UK's private non-financial sector. Figure A.2 in the Appendix plots the raw time series. The series begins in 1999Q4, constraining us to begin the exercise on this date.

Figure 5: *Historical decomposition of GDP-at-Risk for the United Kingdom, 1999-2014*

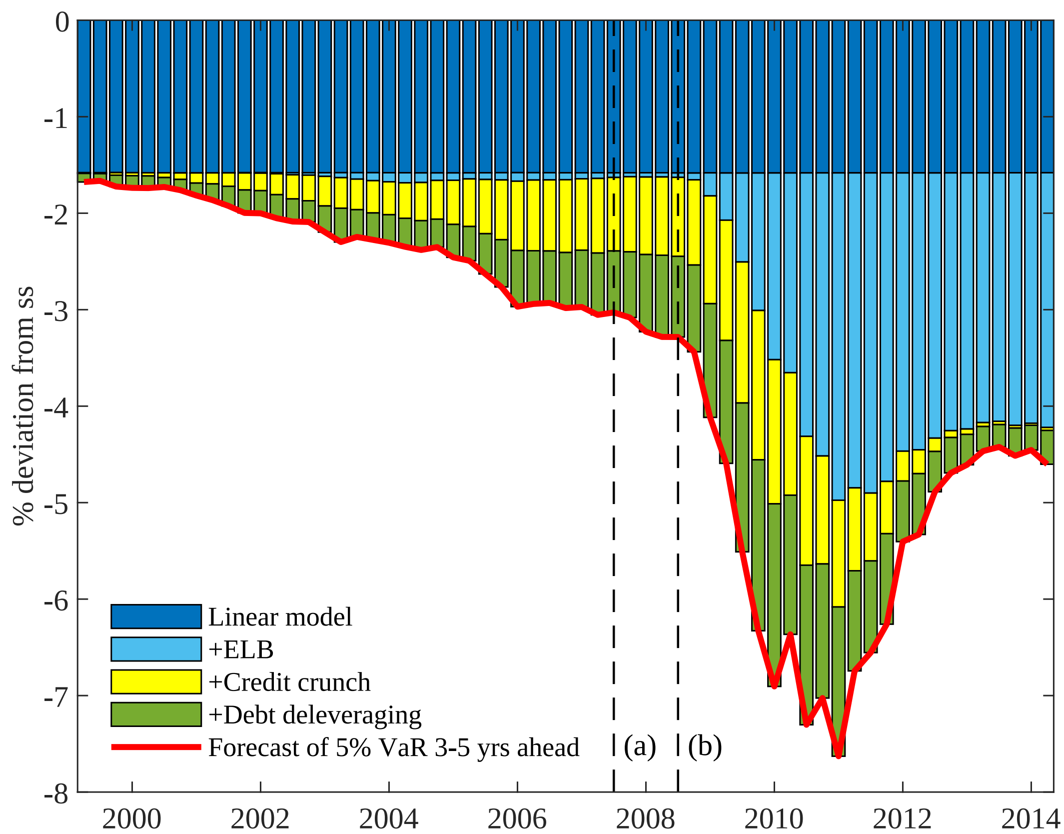

Notes: This figure presents the model's projected 5% level of GDP-at-Risk 3-5 years ahead in the United Kingdom between 1999 and 2014. For each date, we condition the model on data outturns and simulate 20 quarters forward, repeating the process 5000 times. GDP-at-Risk for each date is given by the average 5th percentile of output across quarters 12 to 20, averaged across simulations. The bars show the Shapley value contributions of the three nonlinear constraints alongside the baseline linear model. The dashed vertical lines (a) and (b) denote two key dates during the crisis: Northern Rock's liquidity support from the Bank of England (14 September 2007) and Lehman Brothers' filing for Chapter 11 bankruptcy protection (15 September 2008) respectively.

years of around -3% – almost double the normal risk level. This is below the realised peak decline in GDP of -5.3% in 2009 Q1, however.

After the collapse of Lehman Brothers, Bank Rate was reduced sharply. While this supported the economy in the short-run, it also reduced the scope for the central bank to cushion future shocks. GDP-at-Risk in the medium term therefore deteriorated. Consistent with this, the contribution of the effective lower bound can be seen to grow significantly from late 2008 onwards, leading to a peak risk level of around -7.5% in 2011. Thereafter, the risk outlook began to improve, driven by the long and gradual rebuild of bank capital buffers via the Basel III process, alongside the post-crisis decline in debt service costs. The risk level settled at around -4.5% by mid-2014, with the limited room for further interest rate stimulus being the main cause of higher-than-normal GDP-at-Risk.[12]

These precise risk estimates are unsurprisingly sensitive to the thresholds assumed for the credit crunch and debt deleveraging regimes – thresholds that are unobservable and over which there is significant uncertainty. We explore the impact of alternative plausible thresholds in Figure A.7 and find that our results are broadly robust to these alternative threshold values. The scale of risk estimates is sensitive, however – particularly in the 2009-2012 period.

5.2. Heightened inflation scenario

We use our model to investigate how a shock to inflation affects tail risks. We conduct a scenario analysis in which we feed in a fairly persistent inflation shock that peaks at 8% at the end of 2022 and then remains at 5-6% in 2023-2023, only falling back to steady state by around mid-2026.

Figure 6 plots the tail risk of the GDP forecast in the inflation scenario compared to a baseline where the economy is growing at trend. Overall, high inflation unambiguously worsens tail risks over the next 2-3 years. The model's estimate of the worst-case outcome for the economy is 1.5% points worse in 2023. While around half of this would be captured by standard macroeconomic models (navy blue bars), the remaining effect is due to amplification from the risk of higher interest rates pushing some borrowers' debt burdens into unsustainable territory leading them to deleverage and cut consumption (green bars). Given the high levels of capital in the model, as per calibration for the UK, banks do not amplify the shock further, so that risks from a potential credit crunch are very low (tiny yellow bars). Interestingly, by 2025 GDP at-Risk is back to baseline – and even improved – as these recessionary forces are offset by the benefit of having more monetary policy headroom to cushion other shocks in the future (light blue bars).

6. Conclusion

This paper presents a semi-structural New Keynesian model augmented with three key occasionally binding constraints to generate meaningful tail risks in the GDP distribution: the effective lower bound, household deleveraging, and credit crunch.

We find that all of these non-linearities increase tail risks, but that the ELB causes the strongest amplification of negative shocks. We also find that they interact and amplify each other: For example, an economy prone to debt deleveraging will experience significantly more credit crunch and effective lower bound episodes than otherwise. Moreover, as the effective lower bound becomes more proximate, the frequency of credit crunch episodes increases significantly. This interaction implies that macroprudential policy could be used to provide more monetary policy headroom and vice versa, although this only holds up to a certain threshold.

[12]We choose to end this exercise in 2014 because the Bank of England introduced limits on permissible loan to income ratios from July of that year; assessing the impact of these constraints on the risk level is beyond the scope of our analysis in this paper.

Figure 6: *GDP-at-Risk in a high inflation scneario*

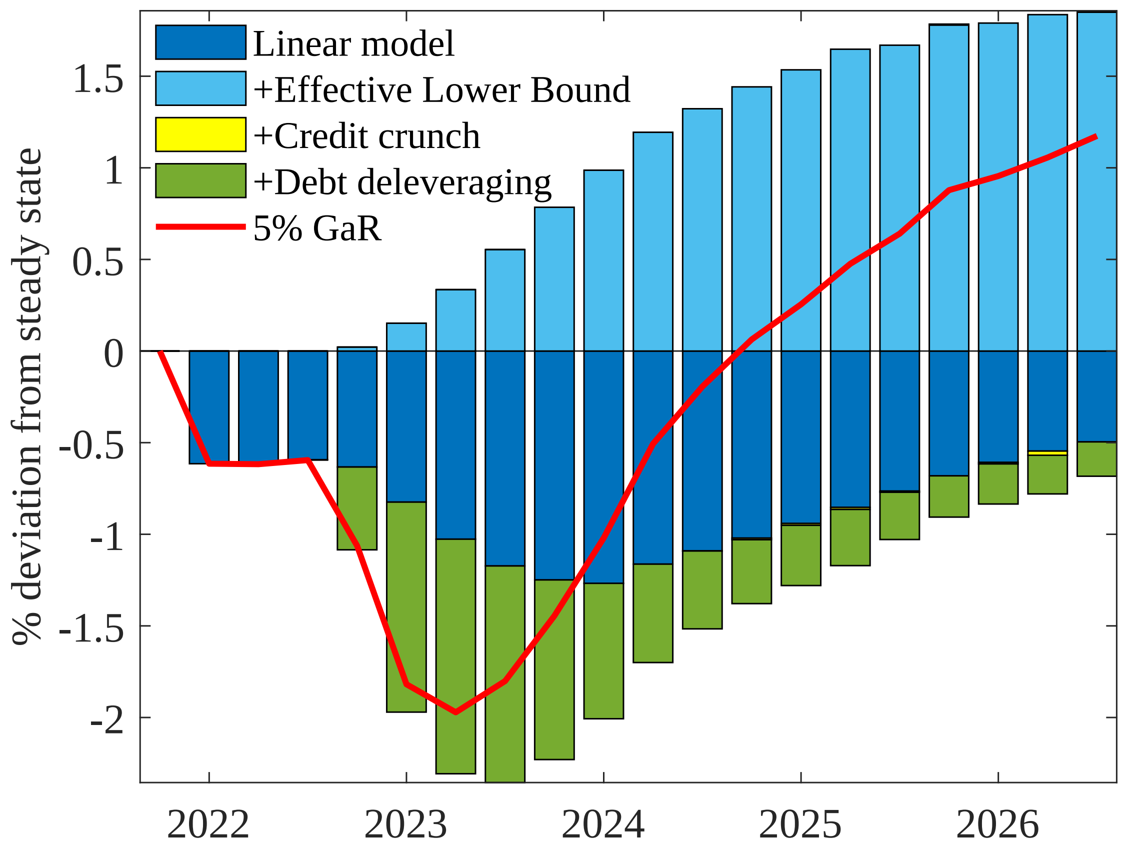

Notes: This figure presents the model's projected 5% level of GDP-at-Risk 3-5 years ahead in the United Kingdom for a high inflation scenario with inflation peaking at 8% at the end of 2022 and then remaining at 5-6% for 2023-2023. The red line indicates the GDP-at-Risk forecast for each date is given by the average 5th percentile of output across quarters 12 to 20, averaged across simulations. The bars show the Shapley value contributions of the three nonlinear constraints alongside the baseline linear model.

We apply the model to the Financial Crisis period and a high inflation scenario to understand the role of the constraints for tail risks in different contexts. Pre-financial crisis the depletion in bank capital buffers and the build-up in private sector nonfinancial debt contributed to the worsening of the growth outlook in the run up to the Global Financial Crisis. In the post-crisis environment, the limited space for monetary policy to cushion adverse shocks was a leading driver for tail risks. In a high inflation world, the model suggests that in the short-term higher interest rates put pressure on debt servicing costs leading to a greater risk of debt deleveraging by heavily indebted households. There is also a risk of higher loan defaults eroding banks' equity capital, which could lead banks to tighten lending conditions – although this effect is small in our model given the size of banks' capital buffers. In the medium term, financial stability risks recede as central banks have greater scope to cut interest rates in a stress, cushioning the effects of negative shocks.

In the ongoing work, we are exploring extending the model along a number of interesting dimensions. One natural area of interest is including a role for countercyclical macroprudential tools such as the CCyB, and whether a strategy of adjusting such tools based on forecasts of the risk level is desirable or not.

References

Adrian, T. and Shin, H., (2010), “Liquidity and leverage” Journal of Financial Intermediation, Vol. 3, Issue 3.

Adrian, T. and Shin, H., (2014), “Procyclical leverage and value-at-risk.” Review of Financial Studies 27 (2): 373–403.

Adrian, T., Boyarchenko, N., and Giannone, D., (2019),”Vulnerable growth” American Economic Review 109(4), April 2019, pp. 1263-89.

Adrian, T. and Duarte, F., (2017), ”Financial vulnerability and monetary policy,” 2017 Meeting Papers 391, Society for Economic Dynamics.

Aikman, D., Bridges, J., Hacioglu Hoke, S., O’Neill, C., and Raja, A., (2019), Credit, capital and crises: a GDP-at-Risk approach”, Bank of England Working Paper No.824.

Alessandri, P. and Nelson, B., (2015), “Simple banking: profitability and the yield curve” Journal of Money Credit and Banking, 47(1), 143-175.

Amihud (2002), “Illiquidity and stock returns: cross-section and time-series effects” Journal of Financial Markets, Volume 5, Issue 1, Pages 31-56.

Aoki, K., Proudman, J., and Vlieghe, G., (2002) ”Houses as collateral: has the link between house prices and consumption in the U.K. changed?,” Economic Policy Review, Federal Reserve Bank of New York, issue May, 163-177.

Bernanke, B. and Blinder, A., (1992), ”The federal funds rate and the channels of monetary transmission” American Economic Review, American Economic Association, vol. 82(4), pages 901-921, September.

Borio, C. and Lowe, P., (2002), “Asset prices, financial and monetary stability: exploring the nexus,” BIS Working Papers no 114.

Brunnermeier, M. K. and Pedersen, L. H., (2009), “Market liquidity and funding liquidity.” Review of Financial Studies 22 (6): 2201–38.

Bunn, P. and Rostom, M., (2015), ”Household debt and spending in the United Kingdom,” Bank of England working papers 554, Bank of England.

Burgess, S., Fernandez-Corugedo, E., Groth, C., Harrison, R., Monti, F., Theodoridis, K., and Waldron, M., (2013), ”The Bank of England’s forecasting platform: COMPASS, MAPS, EASE and the suite of models,” Bank of England working papers 471, Bank of England.

Campbell, J. and Mankiw, G., (1989). “Consumption, income and interest rates: reinterpreting the time series evidence” NBER Macroeconomics Annual 1989, Volume 4.

Carlstrom, C. T., Fuerst, T. S., and Paustian M., (2010). ”Optimal Monetary Policy in a Model with Agency Costs,” Journal of Money, Credit, and Banking, vol. 42, pp. 37-70.

Cecchetti, S., and Suarez, J. (2021). ”On the Stance of Macroprudential Policy”, ESRB: Advisory Scientific Committee Reports 2021/11.

Chavleishvili, S., Fahr, S., Kremer, M., Manganelli, S. and Schwaab, B. (2021). ”A risk management perspective on macroprudential policy,” Working Paper Series 2556, European Central Bank.

Curdia, V., and Woodford, M., (2010), “Credit spreads and monetary policy”, Journal of Money, Credit and Banking, Volume 42, Issue s1, September, 3-35.

Curdia, V., and Woodford, M., (2016), “Credit frictions and optimal monetary policy”, Journal of Monetary Economics, Volume 84, December, 30-65.

Detken, C., Weeken, O., Alessi, L., Bonfim, D., Boucinha, M. M., Castro, C., Frontczak, S., Giordana, G., Giese, J., Jahn, N., Kakes, J., Klaus, B., Lang, J. H., Puzanova, N. and Welz, P. (2014), “Operationalising the countercyclical capital buffer: indicator selection, threshold identification and calibration options”, Occasional Paper Series, No 5, European Systemic Risk Board.

Duca, J.V., Muellbauer, J., Murphy, A., (2011), “Shifting Credit Standards and the Boom and Bust in US House Prices” SERC Discussion Paper 76.

Eggertsson, G. B. and Krugman, P. (2012), “Debt, deleveraging, and the liquidity trap: a Fisher-Minsky-Koo approach” Quarterly Journal of Economics, 127, 1469–1513.

Gertler, M., and Karadi, P., (2011), “A model of unconventional monetary policy” Journal of Monetary Economics, 58(1), 17-34.

Gilchrist, S., and Zakrajsek, E., (2012), “Credit spreads and business cycle fluctuations.” The American Economic Review 102 (4): 1692–720.

Guerrieri, L., Iacoviello, M., Covas, F., Driscoll, J. C., Jahan-Parvar, M., Kiley, M., Queralto, A., and Sim, J. (2019), “Macroeconomic effects of banking-sector losses across structural models” International Journal of Central Banking. Vol. 15, No. 3.

He, Z., and Krishnamurthy, A., (2013) “Intermediary asset pricing” American Economic Review Vol. 103, No. 2, 732-70.

Israelsen, C. (2009). ”Refining the Sharpe Ratio.” Journal of Performance Measurement. Vol. 13(3).

Kiley, M. T., and Roberts, J. M., (2017), ”Monetary policy in a low interest rate world,” Finance and Economics Discussion Series 2017-080, Board of Governors of the Federal Reserve System (US).

Kiyotaki, N., and Moore, J., (1997), “Credit cycles”, Journal of Political Economy 105, 211-248

LLoyd, S., Manuel, E., and Panchev, K. (2021) ”Foreign vulnerabilities, domestic risks: the global drivers of GDP-at-Risk”, Bank of England Working Paper No 940.

Loria, F., Christian M. and Zhang, D. (2019). "Assessing Macroeconomic Tail Risk," Finance and Economics Discussion Series 2019-026, Board of Governors of the Federal Reserve System (U.S.).

Mian, A. and Sufi, A., (2009) "The consequences of mortgage credit expansion: evidence from the U.S. mortgage default crisis" Quarterly Journal of Economics, 124, 1449–1496.

Muellbauer, J. (2012), "When is a housing market overheated enough to threaten stability?", pp. 73-105, in: A. Heath, F. Packer and C. Windsor (eds), Property Markets an d Financial Stability, Proceedings of a Conference Held in Sydney on 20–21 August, Reserve Bank of Australia, Sydney

Plagborg-Moller, M., Reichlin, l., Ricco, G., and Hasenzagl, T. (2020). "When is Growth at Risk?" Brookings Papers on Economic Activity 2020 (Spring): 167-229.

Ravenna, F. and Walsh, C., (2006), "Optimal monetary policy with the cost channel", Journal of Monetary Economics, Vol. 53 (2), March.

Romer, C., D., and Romer, D., H., (2017), "Why some times are different: macroeconomic policy and the aftermath of financial crises," NBER Working Papers 23931, National Bureau of Economic Research, Inc.

Svensson, L. E. O., (2018) "Housing prices, household debt, and macroeconomic risk: problems of macroprudential policy".

Suarez, J. (2021). "Growth-at-risk and macroprudential policy design", ECB Occasional Paper Series No 19.

Woodford, M., (2003), "Optimal interest rate smoothing" The Review of Economic Studies, Volume 70, Issue 4, October 2003, 861–88.

Appendix A: Supplementary charts and tables

Figure A.1: *UK time series*

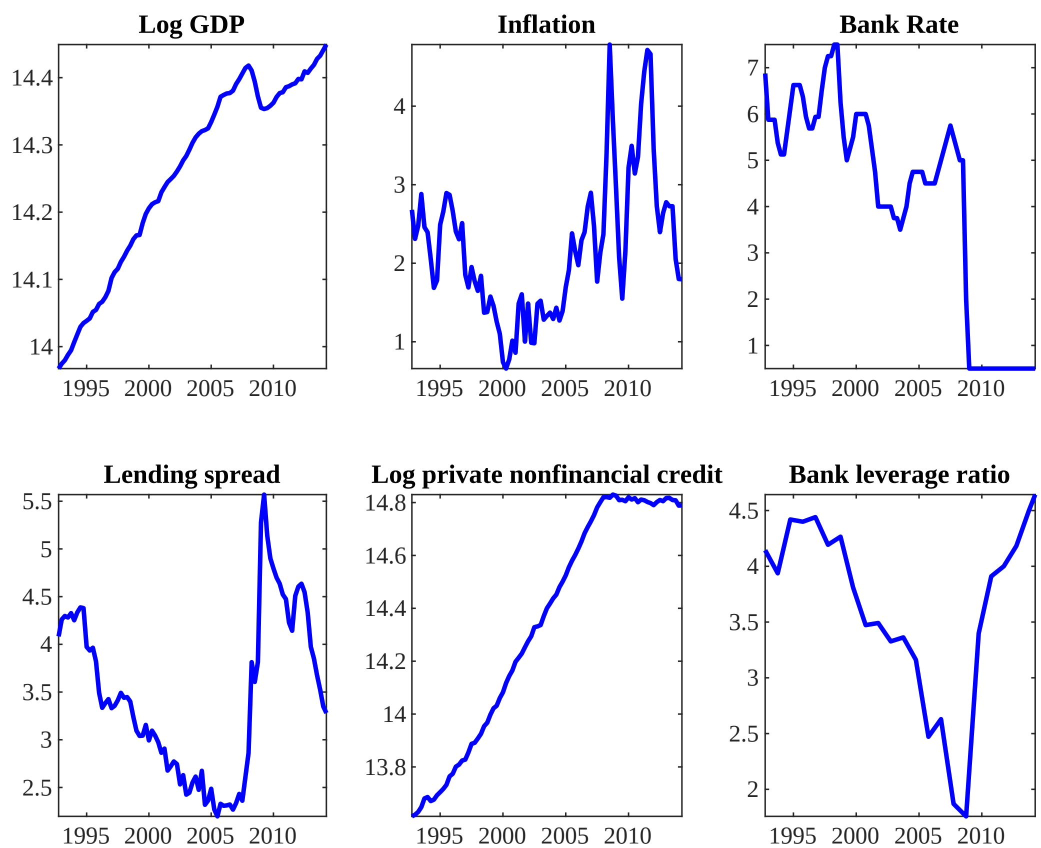

Notes: This figure reports time series of the data series used in the analysis .

Figure A.2: *UK debt service ratio*

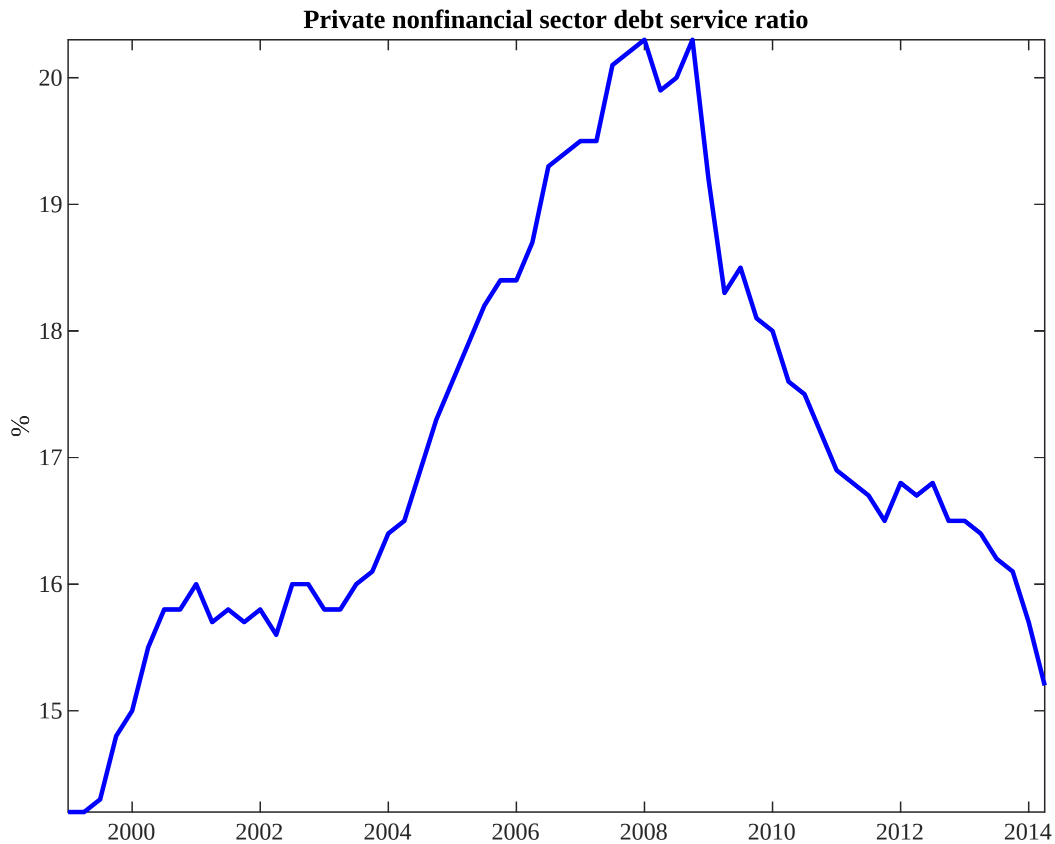

Notes: This figure shows debt service payments relative to income in the UK's private nonfinancial sector. Source: Bank of International Settlements.

We calibrate other key parameters of the model to match certain moments observed in the data and generated by other models in the literature. In particular, we choose the following specific calibration targets:

– First, the responses of output and inflation to an exogenous monetary policy shock in COMPASS (Burgess et al. (2013));

– Second, the response of banks' net interest margin following an exogenous monetary policy shock in Alessandri and Nelson (2015);

– Third, the relationship between GDP, impairments and bank capital in the Bank of England's annual stress tests of the UK banking system (Bank of England (2018));

– Fourth, the relationship between ex ante household debt growth and ex post GDP declines around financial crisis episodes reported in Aikman et al. (2019).

Table A.1 compares the performance of the model to these target moments. In each case, we report peak responses over the relevant horizon. Overall, the fit of the model is reasonable, although there are trade-offs in which moments to hit.

Responses to a 25bps monetary policy shock are close to those in reported in Burgess et al. (2013), with an almost identical output response, but a somewhat smaller impact on inflation (-0.05 % points vs -1.1 % points). To calibrate the relationship between interest rates and banks' net interest margins (as governed by δ_r and δ_s), we match the results reported in Alessandri and Nelson (2015), whereby positive shocks to interest rates initially depress banks' NIMs (as funding reprices more quickly than loans) before eventually exerting a modest positive impact.

To examine the relationship between macro variables and banking system losses and capital, we simulate a macroeconomic stress chosen to match the scenario used by the Bank of England in its 2018 stress test (Bank of England (2018)). In this test, UK output contracts by 4.7%, Bank Rate is assumed to rise to 4%, and investment grade bond spreads increase by 380bps. In our scenario, we apply a combination of IS, Phillips curve and loan spread shocks calibrated to deliver similar responses: our scenario has output contracting by 4.8%, Bank Rate rising to 4.5% and lending spreads rising by almost 340bps. In response to these shocks, the banking system's leverage ratio declines by 1.15% points, closely matching the -1.1% impact on bank capital in the 2018 stress test. However, we find a markedly smaller impact on impairments, which increase by only 0.6% points in our model compared with 1.4% points in the stress test.

Table A.1: *Comparison of the model with target moments*

Description	Target	Model	Source
(a) Monetary transmission mechanism:			
Response of inflation to +25 bps Bank Rate	-0.11pp	-0.05pp	Burgess et al. (2013)
Response of output to +25 bps Bank Rate	-0.14%	-0.13%	Burgess et al. (2013)
(b) Sensitivity of net interest margin (NIM) to interest rates:			
SR response of NIM to +100 bps Bank Rate	-0.04%	-0.05%	Alessandri and Nelson (2015)
LR response of NIM to +100 bps Bank Rate	0.004%	0.008%	Alessandri and Nelson (2015)
(c) Relationship between output, impairments and bank capital:			
Response of impairments in the BoE stress test	1.40pp	0.61pp	Bank of England (2018)
Response of bank capital in the BoE stress test	-1.10pp	-1.15pp	Bank of England (2018)
(d) Response of output to household deleveraging:			
Debt and recession severity	-0.37	-0.29	Aikman et al. (2019)
Debt service and recession severity	-0.29	-0.50	Drehmann and Juselius (2012)

Notes: This table compares the performance of the model with target moments as described in the text.

Table A.2: *Calibration of the model's shock processes*

Parameter	Description	Value
ρ^y	Persistence of aggregate demand shock	0.95
ρ^π	Persistence of cost-push shock	0.8
ρ^r	Persistence of monetary policy shock	0
ρ^s	Persistence of loan spread shock	0.95
ρ^b	Persistence of credit demand shock	0.99
ρ^u	Persistence of write-off shock	0.85
ρ^d	Persistence of debt-deleveraging shock	0.9
σ_y^2	Variance of aggregate demand shock	0.25^2
σ_π^2	Variance of cost-push shock	0.2^2
σ_r^2	Variance of monetary policy shock	0.1^2
σ_s^2	Variance of loan spread shock	0.1^2
σ_b^2	Variance of credit demand shock	0.75
σ_u^2	Variance of write-off shock	0.15^2

Notes: This table presents calibrated values of the model's shock processes.

Figure A.3: *Autocorrelations of endogenous variables, model vs data*

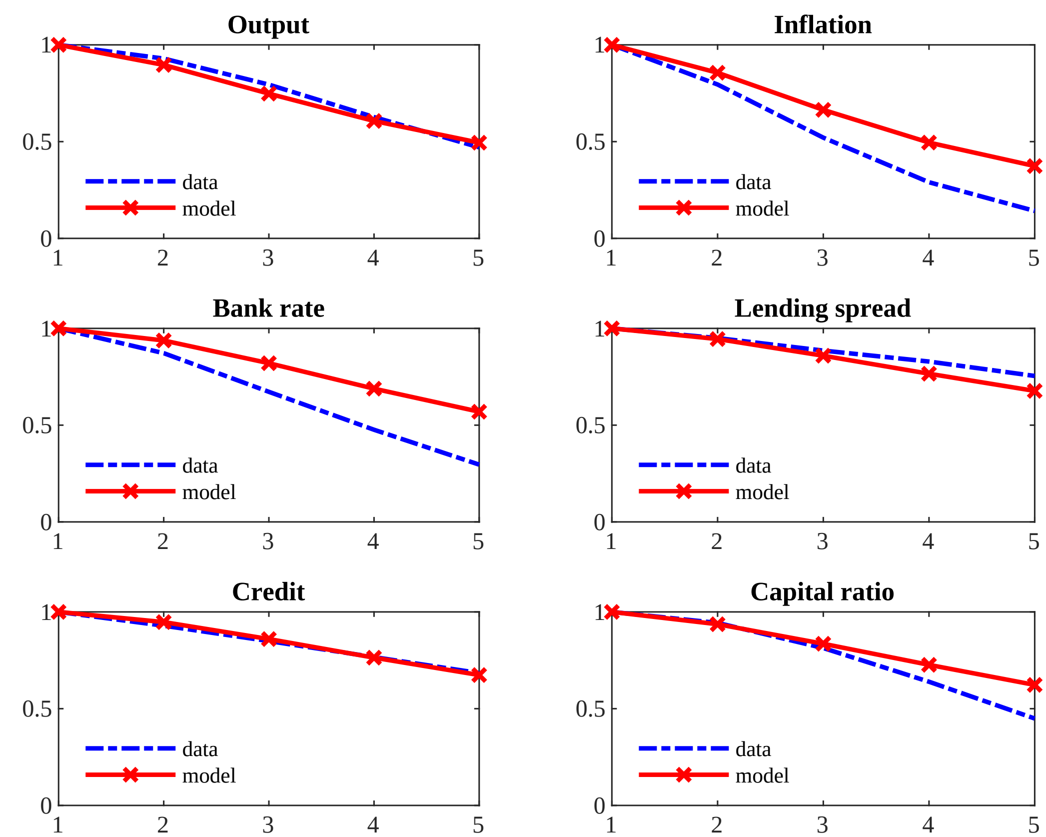

Notes: This figure presents autocorrelation functions in the model versus data.

Figure A.4: *Impulse responses to shocks*

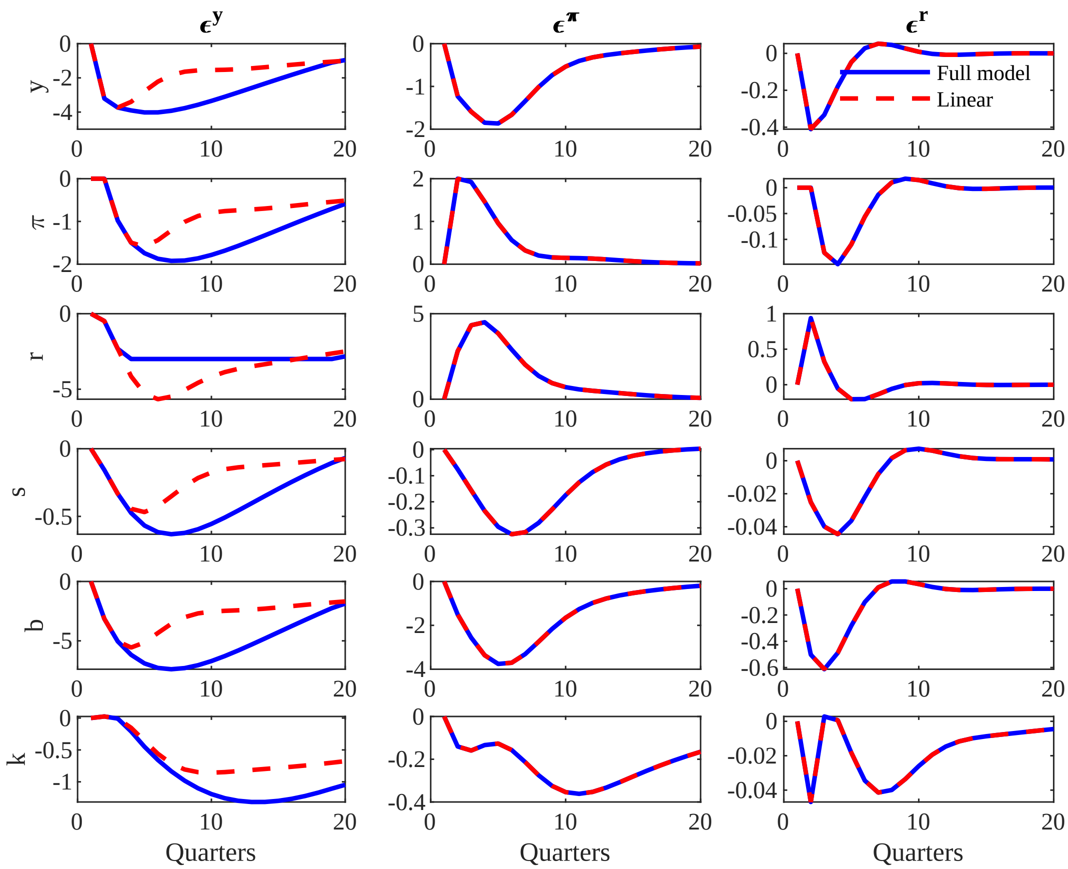

Notes: This figure presents impulse responses of the model's endogenous variables (rows) following shocks to output (ϵ^y), inflation (ϵ^π), and interest rates (ϵ^r) (columns). The blue solid lines show full model responses; red dashed lines show responses in the linear no-constraints model. The y-axis reports % responses for output and credit, and %pt responses for the other variables. The calibration of the steady state is that of the baseline economy, with $\bar{r} = -3$, $\bar{k} = -2$, $\bar{dsr} = 10$.

Figure A.5: *Impulse responses to shocks*

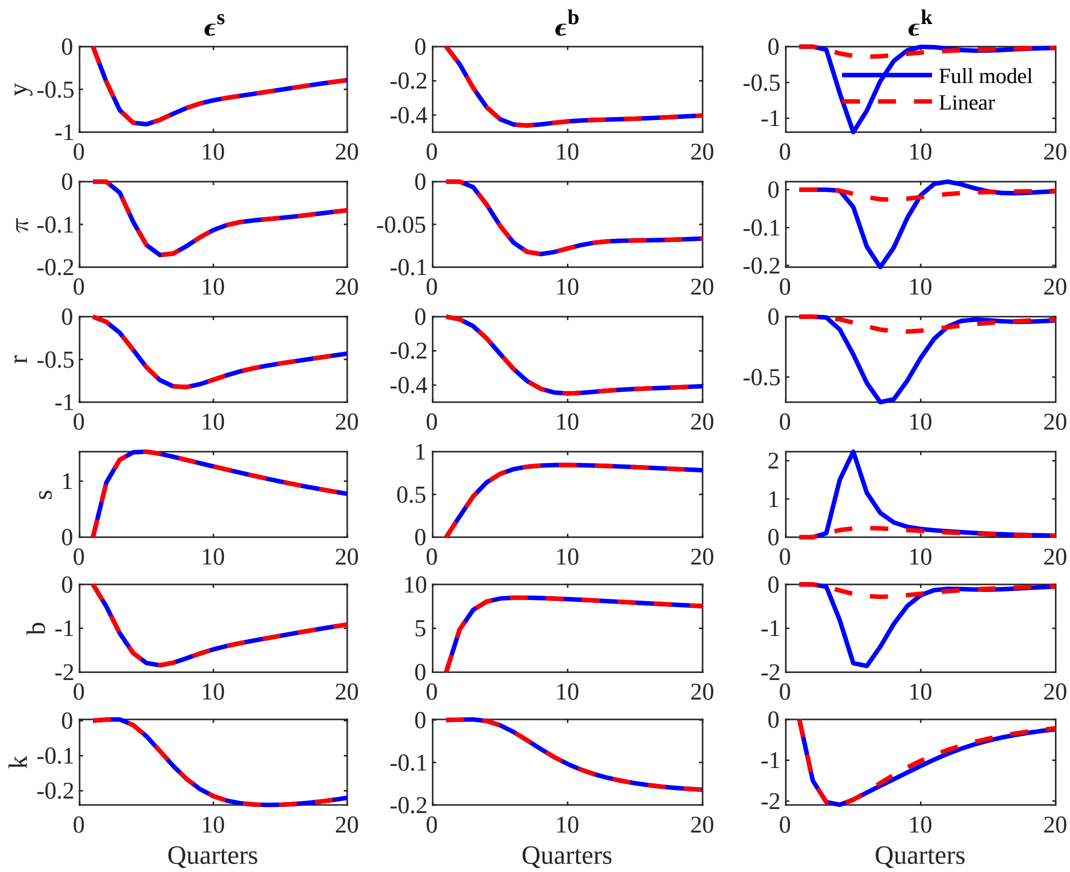

Notes: This figure presents impulse responses of the model's endogenous variables (rows) following shocks to loan spreads (ϵ^s), credit demand (ϵ^b), and bank capital (ϵ^k) (columns). The blue solid lines show full model responses; red dashed lines show responses in the linear no-constraints model. The y-axis reports % responses for output and credit, and %pt responses for the other variables. The calibration of the steady state is that of the baseline economy, with $\bar{r} = -3$, $\bar{k} = -2$, $d\bar{s}r = 10$.

Figure A.6: *Example of dynamics of bank leverage with recapitalisation*

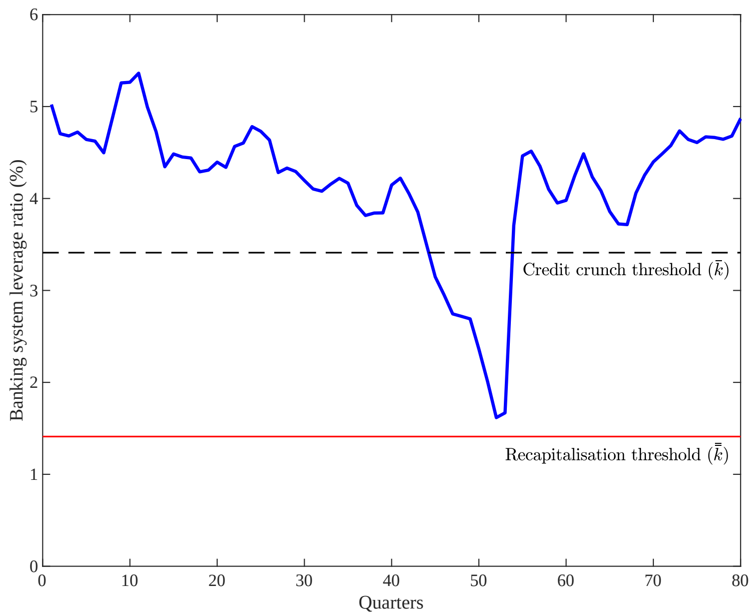

Notes: This figure illustrates the impact of a recapitalisation on the dynamics on bank leverage in one simulation run. We transform the model's raw output, which is expressed as a %pt deviation from steady state, to levels of the leverage ratio by assuming the UK banking system was at steady state in 2017Q4 with a leverage ratio of 5.4%. In the simulation, the banking system suffers mounting credit losses eroding its equity cushion such that its credit crunch threshold is breached. Thereafter, the banking system reduces credit availability sharply, weakening economic activity and increasing credit losses further. This adverse feedback loop is finally halted in period 52 when a recapitalisation occurs, preventing what would otherwise be a breach of the point of non-viability. Once recapitalised, bank credit conditions normalise and the economy begins to recovers.

Figure A.7: *Sensitivity of GDP-at-Risk historical decomposition to alternative constraint thresholds*

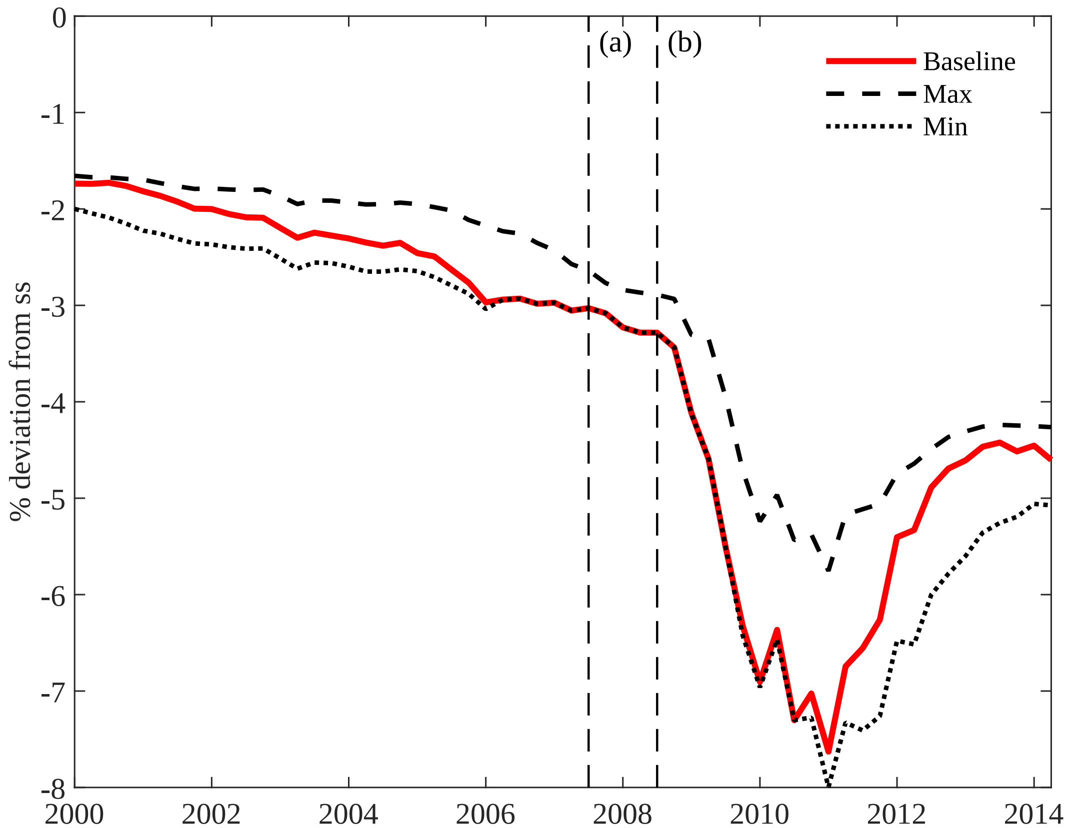

Notes: The red (solid) line is the 5% GDP-at-Risk projection 3-to-5 years ahead conditional on realised data outturns and a given distribution of shocks. See the note to Figure 5. The black dashed and dotted lines are the maximum and minimum GDP-at-Risk values each quarter across 4 alternative simulations with $\bar{k} = 2.3\% \pm 1$ and $\bar{dsr} = 20\% \pm 2$. The dashed vertical lines (a) and (b) denote the dates when Northern Rock sought liquidity support from the Bank of England (14 September 2007) and when Lehman Brothers filed for Chapter 11 bankruptcy protection (15 September 2008) respectively.

Appendix B: Shapley value methodology

In this appendix, we set out details of the Shapley value contributions analysis used in sections 4.1 and 5.1. Implementing this requires that we simulate all combinations in which constraints are applied and assigning each constraint its average contribution across these combinations.

In particular, if we start with the linear model as the baseline, there are $3! = 6$ possible orders in which the nonlinear constraints can be applied. These are: (ELB, CC, DSR), (ELB, DSR, CC), (CC, ELB, DSR), (CC, DSR, ELB), (DSR, ELB, CC), and (DSR, CC, ELB).

The next step is to calculate the marginal impact on the overall risk level of applying each constraint in each of the above orderings. These are obtained by comparing nested versions of the model, which isolate the impact of the constraint in question. For instance, when the constraints are ordered (ELB, CC, DSR), we must first compare the ELB only model with the linear model to obtain the ELB's impact; then the ELB plus credit crunch model to the ELB only model (model 1) to obtain the bank capital constraint's impact; and third, the ELB plus credit crunch plus DSR constraint model to the ELB plus credit crunch model for the DSR's impact.

The table below lists the full set of cases. Shapley value contributions are given by the mean contributions of each constraint, averaging across the six rows. :

Table B.1: *Calculation of Shapley value contributions*

		Marginal contribution of:		
Constraint order	Linear model	ELB	Bank capital constraint	DSR constraint
(ELB, CC, DSR)	Linear	Model (1)-Linear	Model (2)-Model (1)	Model (3)-Model (2)
(ELB, DSR, CC)	Linear	Model (1)-Linear	Model (3)-Model (4)	Model (4)-Model (1)
(CC, ELB, DSR)	Linear	Model (2)-Model (5)	Model (5)-Linear	Model (3)-Model (2)
(CC, DSR, ELB)	Linear	Model (3)-Model (6)	Model (5)-Linear	Model (6)-Model (5)
(DSR, ELB, CC)	Linear	Model (4)-Model (7)	Model (3)-Model (4)	Model (7)-Linear
(DSR, CC, ELB)	Linear	Model (3)-Model (6)	Model (6)-Model (7)	Model (7)-Linear
Shapley value:	Mean	Mean	Mean	Mean

Notes: The models are as follows: Model 1 = ELB only; Model 2 = ELB + Credit crunch; Model 3 = ELB + Credit crunch + DSR; Model 4 = ELB + DSR; Model 5 = Credit crunch only; Model 6 = Credit crunch + DSR; Model 7 = DSR only. The terminology "ELB only" should be taken to mean that all constraints other than the effective lower bound are sufficiently distant that they can be ignored, and so on.