November 6, 2018 version with online appendix (earlier version: BIS Working Papers, no 594)

A quantitative case for leaning against the wind[†]

Andrew Filardo[*] and Phurichai Rungcharoenkitkul[*]

Abstract

Should a monetary authority lean against the build-up of financial imbalances? We study this policy question in an environment in which there are recurring cycles of financial imbalances that develop over time and eventually collapse in a costly manner. The optimal policy reflects the trade-off between the short-run macroeconomic costs of leaning against the wind and the longer-run benefits of stabilising the financial cycle. We model the financial cycle as a nonlinear Markov regime-switching process, calibrate the model to US data and characterise the optimal monetary policy. Leaning systematically over the whole financial cycle is found to outperform policies of "benign neglect" and "late-in-the-cycle" discretionary interventions. This conclusion is robust to a wide range of alternative assumptions and supports an orientation shift in monetary policy frameworks away from narrow price stability to a joint consideration of price and financial stability.

JEL classification: E52, E58, E32, E37

Keywords: monetary policy, financial stability, leaning against the wind, financial cycle, endogenous risk-taking, time-varying transition probability Markov regime-switching model

[†] We would like to thank Roberto Billi, Claudio Borio, Dietrich Domanski, Mathias Drehmann, William English, Hans Genberg, Karsten Gerdrup, Linda Goldberg, Frank Hansen, Andreas Hoffmann, Mikael Juselius, Jonathan Kearns, Anil Kashyap, Enisse Kharroubi, Luiz Pereira, Hyun Song Shin, Lars Svensson, David Vestin, Anders Vredin, Mark Wynne and seminar participants at the 2018 Annual International Journal of Central Banking Research Conference, the Bank of England, the Bank for International Settlements, the Federal Reserve Bank of Dallas, Halle Institute for Economic Research, the Norges Bank and the Sveriges Riksbank for their insightful comments. All errors are ours.

[*] Bank for International Settlements

1 Introduction

The Great Financial Crisis (GFC) rejuvenated the debate concerning whether and how a central bank should trade off price and financial stability. This renewed interest stands in stark contrast to the pre-crisis consensus, which focused narrowly on price stability. Back then, considerations of private sector indebtedness, financial fragility and asset price bubbles were generally thought to be the rightful purview of prudential regulators and supervisors.

For monetary policy, the pre-crisis approach to financial stability is best characterised as "benign neglect". That is, it was believed that central banks should focus solely on macroeconomic developments and largely ignore financial booms. Monetary policy, however, would stand ready to clean up the mess during and after the bust. Experience with equity price booms in past decades provided broad empirical support for this view. When such booms went bust, there was little lasting impact on the economy. Therefore, as the logic went, why depress the economy and hold inflation below target over time if a once-in-a-lifetime crisis were to materialise randomly and the central bank could clean up at little cost?

Faith in this view was shattered by the GFC for several key reasons. First, the cleaning-up approach did not work. The losses brought about by the GFC were large and very persistent. Second, the GFC appeared to be far from a random event, at least in retrospect. The crisis was preceded by a housing market boom which, together with the widespread use of securitisation, exposed the build-up of systemic risks in the financial sector. Once viewed through the lens of a "boom-gone-bust" process, the GFC was not simply a once-in-a-lifetime event that stood out in economic history. The GFC was simply the latest and more extreme episode.

Since the GFC, our understanding of financial crisis dynamics has evolved in several important ways as interest in financial crisis research was reinvigorated. First, several empirical studies that looked back in history found that costly financial downturns were preceded by significant credit and asset price booms (Schularick and Taylor (2012) who build on Borio and Lowe (2002)). Having more comprehensive data on indebtedness at our disposal has opened up the possibility that financial cycles can now be better measured and tracked than in the past. Second, empirical studies have also documented that various financial variables can be used to predict financial crises. Financial imbalances, in particular, have been identified as key predictors of financial busts and subpar economic performance (Borio and Drehmann (2009), Drehmann et al (2012), Jorda et al (2013) and Mian et al (2016)). Third, recent theoretical research has put a spotlight on the role of banks and financial intermediaries in macroeconomic models. This line of research has highlighted various types of financial frictions that amplify economic shocks and exacerbate business cycles (see Brunnermeier Oehmke (2012) for a review).

Of particular note, the endogenous risk-taking channel of monetary policy has taken on greater prominence recently in both empirical and theoretical research. This channel can be activated in several ways. First, easy monetary policy can encourage banks to seek higher returns and take more risks on their loan books (Borio and Zhu (2008), Jimenez et al (2012) and Dell'Ariccia et al (2013)). Banks may also increase their reliance on shorter-term funding to take advantage of lower funding costs (Adrian and Shin (2010)). As a result, this excessive risk-taking behaviour leaves the banking sector more vulnerable to shocks. Second, an easier policy can compress risk premia and push asset prices above levels justified by fundamentals. This raises the risks of an asset price bubble that can correct itself abruptly and damage the balance sheets of investors. In extreme situations, the correction can trigger a fire sale of collateralised assets and intensify an economic downturn. Third, easier monetary policy can accelerate the growth of shadow banking activity and liquidity creation outside of the regulatory umbrella. This increases the fragility of the financial system as higher-risk financial intermediaries are subject to costly runs (Moreira and Savov (2017)).

Our evolving understanding of financial crisis dynamics has given rise to several perspectives on the debate. On one side, some have called for greater reliance on macroprudential tools. Under this view, monetary policy should take a backseat. Others have suggested that macroprudential tools should be used as the first line of defence and monetary policy only as a last resort.

However, there is uncertainty about the effectiveness of macroprudential tools, not least owing to the fact that risk-taking channels appear to be so diverse and not fully understood yet. This uncertainty suggests that the targeted approach of macroprudential tools may at best be an incomplete safeguard. And, over time, macroprudential tools can potentially be circumvented via regulatory arbitrage and creative financial engineering.

The shortcomings of macroprudential tools have left open an important role for monetary policy to lean against the wind. The potential benefit from the use of monetary policy is that the policy rate is not only a powerful macroeconomic tool but also the universal price of leverage and risk. In this view, macroprudential tools and the policy rate should be seen as complements, not as substitutes. Taken together, the limits of macroprudential tools and the potential of the policy rate to influence the financial cycle provide a *prima facie* case for monetary policy to be used in the pursuit of financial stability.

However, the feasibility of using the policy rate does not guarantee its desirability. A number of recent studies have begun exploring the costs and benefits of using monetary policy in the pursuit of financial stability. The emphasis has been on quantifying the costs and benefits with explicit models. One influential strand of recent research has been pioneered by Svensson (2017), Riksbank (2013) and IMF (2015). That strand considers explicit cost-benefit calculations in monetary policy models that feature financial crises. It identifies three key considerations when calibrating the costs and benefits: (i) how much leaning is needed to curb credit growth (those studies' measure of financial imbalances)?; (ii) how do changes in credit growth affect the likelihood of a future financial crisis?; and (iii) how costly is pre-emptive policy in terms of short-term macroeconomic costs (for example, on unemployment and output)? The authors of those studies find evidence against leaning and argue that this is a robust finding.

This no-leaning prescription has been challenged, however. Adrian and Liang (2018), for example, find that the result is not robust to the relaxation of some key assumptions. In particular, if leaning helps lower the eventual cost of a crisis sufficiently, then it is beneficial. This alternative assumption is consistent with the historical evidence that strong financial imbalances tend to intensify economic downturns (Jorda et al (2013) and Juselius et al (2017)).

More fundamentally, the benefit of leaning against the wind may be better appreciated if one recognises (i) the endogenous process governing the slow build-up of financial imbalances, which culminates in a crisis if there is sufficient momentum, and (ii) the systematic influence of policy over the entire financial cycle. Indeed, one could argue that leaning is less about averting an imminent crisis and more about fostering financial stability at all times even when a crisis remains a very remote possibility (Borio and Lowe (2002), BIS (2016) and Juselius et al (2017)). In this case, a leaning policy might be better thought of as an integral part of the monetary policy *framework*, rather than as an occasional deviation from a conventional inflation-targeting approach. This paper attempts to address this gap in the literature.

The paper proposes a dynamic model for evaluating leaning-against-the-wind policies in the presence of recurring financial cycles. The model consists of the conventional macroeconomic block, augmented with a financial cycle block that describes the evolution of financial imbalances and their impact on the economy. In this paper, the financial cycle is persistent and endogenous. In a boom phase, financial imbalances grow over time. Once they reach a sufficiently high level, there is a progressively rising probability that the economy will switch into a financial downturn phase, during which imbalances shrink.[2] The pace at which financial imbalances build up is influenced by a leaning policy, not least reflecting the basic features of the risk-taking channel of monetary policy. Monetary policy can then play a role in constraining the accumulation of imbalances, and consequently lessen the duration (and thus the total cost) of a crisis.[3]

A number of other studies have also used regime-switching models to introduce a financial crisis module to monetary policy modelling. Like ours, these studies posit that the probability of a crisis event depends on some financial variable (leverage in Woodford (2012), credit growth in Ajello et al (2015), the debt level in Alpanda and Ueberfeldt (2016), and credit in a model of a small, open economy in Gerdrup et al (2017)).

Our paper distinguishes itself from these prior studies by considering a persistent dynamic process governing the evolution of the financial imbalance variable consistent with the data. These studies assume that any unexpected opening up of financial

[2] Our characterisation of the financial cycle is consistent with the recent theoretical literature. It highlights a variety of market imperfections that could give rise to a persistent build-up of financial imbalances, followed by highly nonlinear adjustments (for a review, see Brunnermeier and Oehmke (2012)).

[3] The unavoidable cost of a crisis captures the cyclical fall in output that cannot be offset by easing policy. This may be due to financial frictions during downturns or effective lower bound on the policy rate. The argument for leaning would be strengthened if a bust led to a stronger and longer-lasting disruption. Gourio et al (2018), for example, allow a crisis to permanently impact technology, thereby strengthening the case for leaning.

imbalances naturally unwind without any policy intervention, so that the risks of a crisis are always expected to diminish over time. It is only when an unlikely sequence of consecutive shocks pushes financial imbalances to a level sufficiently high that a crisis becomes possible and a leaning policy might be justified. In our model, the financial cycle exhibits strong persistence from both *ex post* and *ex ante* perspectives. We show that, in this environment, there is a role for monetary policy to react early on and pre-empt large imbalances from building up in the first place. The optimal policy prescribes systematic leaning rather than occasional tightening when crisis risks are deemed to be high. In our model, crises may be unlikely and infrequent, but they are not random events.

The rest of the paper is organised as follows. The next section lays out the dynamic monetary policy model featuring costly financial cycles and characterises the optimal monetary policy. Section 3 explores the robustness of the findings under alternative sets of assumptions. The final section concludes that there is a strong case for policy frameworks that systematically lean against the wind. The main benefit of such policies lies not in averting an impending crisis but in lessening the likelihood and severity of financial cycles over time.

2 Optimal policy under an endogenous financial cycle

We construct a model comprised of three key building blocks: a macroeconomic block, a financial cycle block and a monetary policy block. We calibrate the dynamic model to US data and characterise the optimal degree of leaning.

The financial cycle block plays a central role and is modelled to capture several key empirical features. First, financial cycles go through periods of booms and busts. Severe financial cycles result in very costly crises while less extreme swings in financial conditions impose a lower but non-zero cost. Second, the transition from a financial cycle boom to bust tends to be sharp and non-linear, with the size of the bust increasing in the degree of financial imbalances that need to be unwound. We capture such dynamics via a 2-regime Markov process with a boom regime and a bust regime. Furthermore, the transition probabilities are modelled as being time-varying, depending on the degree of financial imbalances and implying an endogenously persistent financial cycle. It is important to note that an endogenously persistent cycle is very different from a conventional persistent-shock process. The endogenous cycle is one in which the up and down swings in the financial cycle can occur without reinforcing shocks. A conventional persistent-shock model in contrast relies on sustained, correlated shocks to generate financial ups and downs. Without shocks, this type of process is strongly self-correcting. Third, financial crises can be costly and the costs cannot be completely offset by monetary policy. One can think of this as a shorthand for the assumption that monetary policy faces limits in the ability to clean-up during a financial deleveraging, eg due to the balance sheet effects of debt overhang and persistent effects on confidence. This type of unavoidable loss creates incentives to try to prevent large financial imbalances from occurring in the first place.

2.1 Model specification

2.1.1 Macroeconomic block

The macroeconomic block is an IS curve describing the behaviour of the output gap, y_t, as a function of the real policy rate gap, r_{t-1}. For simplicity, it is assumed to take a backward-looking form as in Rudebusch and Svensson (1999)

$$y_t = \beta y_{t-1} + \phi r_{t-1} + e_t^y \tag{1}$$

with $0 < \beta < 1$, $\phi < 0$, and $e_t^y \sim N(0, \sigma_y^2)$. The lagged term captures output persistence. A forward-looking IS curve may be less appropriate in the current setting, as it would imply a spending reduction in anticipation of a crisis, a feature at odds with anecdotal observations that exuberance tends to precede crises (see Ajello et al (2016) for the motivation of a similar assumption).

We abstract from inflation considerations in this baseline specification for simplicity, as our focus is not on the policy trade-offs arising from supply-side shocks. The central bank in this formulation already has an implicit long-run inflation objective, since it aims to stabilise output deviations around the trend, through the business and financial cycles. In Section 3.2, we model inflation more formally with an explicit Philips curve, and show that this extension does not materially change our main findings.

2.1.2 *Financial cycle block*

The financial cycle f_t follows a random walk process with a drift that depends on the regime $s_t \in \{1,2\}$, corresponding to boom and bust phases respectively. The policy rate gap r_t is assumed to influence the financial cycle, with a sensitivity that is regime-dependent and takes the form

$$f_t = f_{t-1} + \alpha_{s_t} + \gamma_{s_t} r_{t-1} + e_t^f \tag{2}$$

where $\alpha_1 > 0$ and $\alpha_2 < 0$ by definition of booms and busts and $e_t^f \sim N(0, \sigma_f^2)$. We assume that $\gamma_1, \gamma_2 \leq 0$, consistent with a risk-taking channel.

The regime s_t follows a (non-homogeneous) Markov process with logistic transition probabilities,

$$P_t(s_{t+1} = 2 | s_t = 1) = \zeta_1 \frac{\exp(\psi_1(f_t - f^H))}{1 + \exp(\psi_1(f_t - f^H))} \tag{3}$$

$$P_t(s_{t+1} = 1 | s_t = 2) = \zeta_2 \frac{\exp(\psi_2(f_t - f^L))}{1 + \exp(\psi_2(f_t - f^L))} \tag{4}$$

The 'sensitivity' parameters $\psi_1 > 0$ and $\psi_2 < 0$ capture the sensitivity of the transition probabilities to developments in the financial cycle. The sign restrictions are such that a probability of transitioning from a boom to a bust grows as financial imbalances build, and the probability of transitioning from a bust to a recovery rises the lower the imbalances. These assumptions capture the endogenous financial cycle idea: a high degree of financial imbalances is a pre-condition for a crisis, a recovery from which is only possible once deleveraging has reduced the stock of imbalances sufficiently.

Around the 'tipping points' $f^H > 0$ and $f^L < 0$, the transition probabilities are most sensitive to changes in f_t. The parameters ζ_1 and ζ_2, with $0 < \zeta_1, \zeta_2 < 1$, scale the transition probabilities in a multiplicative fashion, and are the limiting transition probabilities as $f_t \to \infty$ and $-\infty$ respectively. In sum, eleven parameters define the dynamics of the financial cycle block.

Monetary policy plays a role in stabilising the financial cycle, by counteracting the upward momentum of f_t in booms. And by making large imbalances unlikely to occur on average, leaning also helps limit the expected duration of a bust. In our model, monetary policy can lessen the likelihood and average cost of a downturn.

2.1.3 Central bank objective function

The IS equation describes the normal part of output under the influence of monetary policy. The central bank's objective function is to set monetary policy r_t to maximise the objective function

$$U = E \sum_{t=1}^{\infty} \delta^{t-1} \left(-\frac{1}{2} y_t^2 - cI(s_t = 2) \right) \quad (5)$$

where δ is the discount factor and $I(s = 2)$ is an indicator function equal to one when $s = 2$ and zero otherwise. During the bust phase, we assume the central bank incurs an exogenous output loss $c > 0$ that cannot be offset or mitigated by easier monetary policy. The optimisation is subject to the constraints posed by the IS curve (Equation 1) and the financial cycle dynamics (Equations (2)-(4)).

The assumption that monetary policy easing cannot absorb the bust cost c is crucial. If monetary policy were effective in 'cleaning up after' and capable of offsetting any output loss associated with a crisis, then trivially there would be no gain from leaning to prevent a bust.[4] In the context of the leaning debate, we are ruling out benign neglect monetary policies as being feasible. We do consider an extension where the bust cost $c > 0$ is endogenous. This allows us to address the possibility that a weaker economy is more heavily hit by a bust than a stronger economy. We will show that, as long as there exists a lower bound to the cost c (ie monetary policy cannot fully offset the impact of a financial bust), the optimality of a leaning policy carries over.

2.2 Calibrating the model

2.2.1 Parameterising the macro block and the crisis cost

We assume the following simple IS equation

$$y_t = 0.9 y_{t-1} - 0.4 r_{t-1} + e_t^y, \quad \sigma_t^y = 1.5 \tag{6}$$

The marginal impact of monetary policy on output of –0.4 is calibrated to roughly match the maximum effect used in the FRB/US model. This parameterisation implies

[4] For example, the loss function used in Svensson (2017) has the property that a policy easing can increase output (pre-crisis or during a crisis) and lower the impact of the crisis on welfare. See Section 3.1.5 below for the implications of such a loss function in our setting.

that closing a 1 percent output gap requires a policy rate gap of 2.25 percent. Later, we will consider IS curve specifications with a lower sensitivity to policy rates.[5]

The cost of financial downturn is set to –10 per quarter as the baseline case. In the objective function in Equation (5), this is equivalent to an output gap of roughly –4.5 percent.[6] Finally, the discount factor is set at 0.99. In Section 3, we will consider how variations in all these parameters influence the optimal monetary policy.

2.2.2 *Estimating the financial cycle block*

We estimate the parameters of the financial cycle block using a maximum likelihood procedure for regime-switching models, a la Hamilton (2008), but generalised to allow time-varying transition probabilities (also see Filardo (1994)). The financial cycle f_t is the observed variable that provides inferences about the latent regime s_t. See an online appendix for the details of estimation procedure. The financial cycle block is calibrated using US quarterly data from 1960 Q1 to 2015 Q1 which is an updated series based on the methodology of Drehmann et al (2012). Figure 1 plots the resulting financial cycle series f_t.

5 A simple OLS estimation of Equation 1 yields $y_t = 0.96y_{t-1}0.05r_{t-1}$ (with significant coefficients) from 1954 Q3 to 2015 Q3. The output gap is from the US Congressional Budget Office, and the interest rate gap is the federal funds rate minus the sample average.

6 This assumption is somewhat more conservative than in Ajello et al (2016) which assumes 10 percent output loss for 8 quarters (implying a crisis cost of –50 per period for a quadratic loss function). Assuming that the same cost is incurred for 6 years, which is the expected duration of a bust regime based on our transition probability estimates, their assumption is equivalent to 6.9 percent output loss per period in discounted terms. Our lower cost assumption may be justified by the fact that it is meant to capture the cost of an average financial downturn, including both full-blown crises and less severe deleveraging episodes. It might be thought of as an expected cost of a financial downturn.

[Figure 1 here]

The model's flexibility allows a good fit to the financial cycle data, but raises potential identification issues. The amplitude, duration and turning points of the financial cycle are all time-varying, and observing the time-series f_t alone may not be sufficient to jointly identify all the transition probability parameters precisely. For example, the observed distribution of the cycle's turning points (peaks and troughs) could be explained by different combinations of scale and sensitivity parameters. To address this identification issue, we consider three alternative parameterisations and investigate the robustness of optimal monetary policy. In parameterisation scheme A, we calibrate the tipping points f^H and f^L to match historical averages of the financial cycle peaks and troughs in the United States. In parameterisation scheme B, we place bounds on the sensitivity parameters, $\psi_1 \leq 10$ and $\psi_2 \geq 10$, and estimate all other parameters. This restriction implies that it takes some time for the transition probability to pick up over the cycle. Finally in parameterisation scheme C, we assume bounds $\psi_1 \leq 10$ and $\psi_2 \geq 10$, and impose constraints on the expected turning points to match historical averages.[7]

[Table 1 here]

[7] This is implemented by maximising a weighted sum of a standard likelihood function with a penalty for deviations of the implied turning points from the historical averages (the weight on the latter is set arbitrarily high at 100).

2.2.3 *Estimation results*

The estimates are reported in Table 1. Estimates of all the 'drift' parameters in Equations (2) and σ_f are robust to different estimation schemes. On average, f_t expands slowly during booms, and falls more quickly during busts. Thus, booms tend to last longer than busts on average, in line with what was documented by Drehmann et al (2012) and Claessens et al (2012) for moderate cycles.

Tighter monetary policy helps rein in the boom - a one percentage point increase in the policy rate gap lowers the expected increase in financial imbalances by about 20 percent per quarter (–0.0087/0.0424). Interestingly, a more accommodative monetary policy does not help slow down contractions associated with the bust according to our estimates, so we set the impact of policy during busts to zero. The impact of monetary policy on the financial cycles is therefore restricted to booms. The results are also robust for many other estimation schemes (not shown).[8]

[Figure 2 here]

All three estimation schemes enjoy similarly good fits to the data, but imply different shapes of transition probabilities as shown in Figure 2. In scheme A, the transition probabilities change relatively slowly with the stage of the financial cycle, f, and the model allows for time-varying peaks and troughs by admitting higher

[8] We explore the possibility of nonlinear drifts and policy effect on the financial cycle and confirm that linearity is a good approximation. We also investigate a special case of the policy effect being subject to diminishing return, so even extreme policy intervention cannot eliminate cycles. This special case is rejected by the data. In our estimated model, it is feasible for the central bank to stabilise the financial cycle.

unconditional switching probabilities due to the scale parameters. The probability of entering a bust approaches the maximum of 0.21 per quarter as imbalances become extreme (the probability of exiting a bust approaches 0.26 per period). In schemes B and C, the bounds on the sensitivity parameters ψ are binding, and the transition probabilities are more sensitive to f around the tipping points, f^H and f^L. The limiting transition probabilities in these cases are lower at around 10 percent per period.

Are the restrictions on ψ too stringent? Simple calculations suggest not. With $\psi = 10$, f needs to move by about 0.45 around the tipping points for the transition probability to increase from a near-minimal level of $0.1 \times \zeta$ to a near-maximal level of $0.9 \times \zeta$. In a boom, f rises on average by $\alpha_1 = 0.04$ per quarter when the rate gap is zero, implying that the probability of entering a bust would grow from nearly zero to the maximum level in about 10 quarters. In a bust, the associated length is only four quarters. Note that Figure 1 includes our estimates of the filtered probability of being in a boom (all schemes result in almost identical measures of fit and regime probabilities). The estimated model can explain the United States financial cycle well, and the estimated regime probabilities are broadly in line with how one might visually identify boom and bust episodes, including the latest global financial crisis and the recovery. The model interprets smaller swings in f during 1970s and 1980s as driven by shocks, and not boom-bust cycles.

2.3 Optimal monetary policy

Given the Markovian structure, the monetary policy problem can be expressed in the dynamic programming form where the value function is determined by the state variables (f, y, s):

$$V(f_t, y_t, s_t) = \max_{r(f_t, y_t, s_t)} \left\{ \left(-\frac{1}{2} y_t^2 - cI(s_t = 2) \right) + \delta E_t V(f_{t+1}, y_{t+1}, s_{t+1}) \right\} \quad (7)$$

subject to the estimated laws of motion for f_t, y_t, and s_t. The expectations at time t on the right hand side of Equation (7) are taken with respect to the joint probability distribution of these state variables.

The optimal policy function $r(f, y, s)$ and associated value function $V(f, y, s)$ are numerically computed using the collocation method, which entails solving Equation (7) for a chosen set of discrete coordinates over the state space. We approximate the value function using a cubic spline basis function, construct quadratures to capture shocks, and choose a relatively fine grid along the f dimension to ensure a good approximation to a potentially highly nonlinear optimal policy.[9]

Figure 3 shows the 3-dimensional plots of the results, corresponding to estimation scheme A (for results on other estimation schemes, see the online appendix). A common feature across all schemes is that the boom values $V(f, y, 1)$ strictly exceed the bust values $V(f, y, 2)$ for any state f and y – the two value surfaces are shown in

[9] See Judd (1999) and Miranda and Fackler (2002) for details on the numerical optimisation algorithm.

the top left panel, while the difference between the two surfaces is plotted in the top right panels. The difference is uniformly positive as long as there is a non-zero downturn cost. A consequence of this positive gap is that $V(f, y, s)$ always decreases with f for any y and s. This is because a higher f during booms weakly raises the probability of entering a downturn within any horizon, which reduces the value function, while in busts, a lower f raises the likelihood of improving the value function associated with a recovery. Along the output gap axis, the value function is concave with maxima at $y = 0$.

[Figure 3 here]

The optimal monetary policy during booms and busts is shown in the lower two panels of Figure 3. For a given regime s and the financial cycle f, the optimal policy increases linearly with the output gap, as is standard for a quadratic loss function (lower left panel). But the optimal interest rate varies with f for any given level of the output gap. To see this clearly, we compute the difference between this optimal policy rate and the optimal policy obtained under the assumption of no downturn cost where the central bank simply sets the policy rate to target $y = 0$. In all estimation schemes, it is optimal to set the policy rate gap higher than otherwise during booms – ie leaning is optimal. This is so even at low levels of f when the likelihood of a downturn is extremely remote. The degree of leaning rises with f as the probability of a bust increases, but tapers off as the degree of imbalances reaches an extreme level.

In Figure 4, we show the optimal policy when the output gap is fixed at zero for the three estimation schemes (effectively a cross-sectional slice of the lower-right panel in Figures 3). Here, a policymaker who ignores the financial cycle would have set the policy rate gap uniformly at zero given that the output gap is already closed. The left panel of Figure 4 shows optimal policy during booms, which involves tightening at all stages of the financial cycle. At lower levels of financial imbalances, less leaning is required. But as imbalances build and f approaches the tipping point, it is optimal to lean progressively more. The nonlinear policy response is due to the endogenous regime-switching probabilities. In schemes B and C, the transition probabilities rise rapidly around the tipping points, and hence the optimal policy is to respond more aggressively in these states. The lower tipping points in schemes B and C also suggest that policy needs to be tightened earlier at lower values of f.

[Figure 4 here]

The result supporting the case for early leaning is robust to alternative estimation schemes. At lower levels of financial imbalances, the size of optimal leaning policy is very similar across all estimation schemes – just below 20 basis points at $f = -3$, and about 25 basis points on average for $f \leq 0$. The maximum amount of leaning is between 70–80 basis points, which is optimal as financial imbalances near the tipping points. But even there, an 80 basis points tightening of policy does not appear unusually large – it implies 0.8 x 0.4=0.32 percent lower output. As policy leans systematically throughout the financial cycle, the results indicate that the amount of leaning in any given quarter need not be large.

One interesting feature of the optimal policy function is its bell-shape around tipping points. As f exceeds a tipping point, the transition probabilities become progressively less sensitive to f as the probabilities asymptote to 1. In this case, monetary policy becomes less effective in pre-emptively preventing a bust. As a consequence, it is optimal to lower the degree of leaning at extreme levels of f.

The estimates for γ suggest that, during busts, monetary policy has little impact on the financial cycle dynamics. Thus when the economy just enters a bust and large imbalances remain, an optimising policymaker focuses primarily on macro stability and targets a zero output gap (hence implementing a zero rate gap if the output gap is already zero; right-hand panel of Figure 4). As the size of financial imbalances declines, the transition back to a boom regime becomes more likely and the central bank eases policy. Easing is not intended to influence the financial cycle in this bust state, but to stabilise output in anticipation of future leaning.

3 Extensions

Three types of extensions are explored in this section: alternative parameterisations, the addition of a Phillips curve equation in order to formally introduce inflation into the model, and an assessment of a (short-run) discretionary cost-benefit method of leaning advocated in the extant literature.

3.1 Alternative parameterisation

We consider various scenarios that could affect the optimal degree of leaning, including alternative parameterisations for (i) downturn costs c, (ii) IS curve parameters, (iii) the degree of a central banker's patience, (iv) the persistence of the financial cycle, f and (v) endogeneity of downturn costs with respect to output.

3.1.1 Downturn costs

The implications of different exogenous downturn costs are straightforward as shown Figure 5, where we consider the alternatives $c = 5$ and $c = 20$. We find that the optimal policy is proportional to the downturn costs at all stages of the financial cycle – when the costs double, the optimal degree of leaning is twice as high, and so on. This result generalises the principle that some leaning is optimal as long as the costs of a bust are positive.

[Figure 5 here]

3.1.2 IS curve parameters

Greater output persistence implies more long-lasting effects of leaning on output, raising the marginal cost of leaning. To shed some light on the quantitative implications of this dependence on the extent to which a central bank would want to lean, we compare the optimal policy rules under three alternative IS curve specifications:

Lower persistence: $$y_t = 0.7y_{t-1} - 0.4r_{t-1} + e_t^y \tag{8}$$

Flatter IS curve: $$y_t = 0.9y_{t-1} - 0.1r_{t-1} + e_t^y \quad (9)$$

Low output impact: $$y_t = 0.7y_{t-1} - 0.1r_{t-1} + e_t^y \quad (10)$$

In the first alternative, output is less persistent ('lower persistence' case) than the baseline IS curve. In the second, output is less sensitive to the policy rate ('flatter IS' case). In the third, output has lower persistence and is less sensitive to the policy rate.

Optimal policies under the alternatives are shown in Figure 6 (using estimation scheme A). Lower output persistence indeed leads the central bank to lean more heavily on the financial cycle at every stage, including earlier in the cycle. The optimal degree of leaning is about twice as large as in the baseline.

[Figure 6 here]

But the flatter IS curve has a significantly larger quantitative impact on optimal policy. The impact is many times larger than that coming from a reduction in the policy sensitivity parameter. Combining lower persistence with a flatter slope raises the optimal degree of leaning even further. In other words, leaning yields benefits that grows nonlinearly when IS curves are flatter and less sensitive to the policy rate.

3.1.3 Degree of central bank patience

A greater weight placed on future outcomes should reinforce the case for leaning, as a central bank would at the margin have an incentive to sacrifice more today to avoid future downturns. We evaluate this argument in the context of our model by considering alternative discount factors.

In Figure 7, the baseline optimal policy with the discount factor $\delta = 0.99$ is plotted together with the optimal policy under $\delta = 0.95$ and $\delta = 0.999$ (all based on scheme A and original IS curve). Greater patience significantly strengthens the case for front-loading leaning in the cycle. For example when $\delta = 0.999$, the rate gap is set to about 35 basis points in a boom, even at very low levels of financial imbalances. Moreover, the central bank also leans more forcefully at very high f, even if the marginal influence on the transition probability may be declining. This is because a more forward-looking central bank internalises more fully the endogenous downturn duration. Allowing f to increase implies a longer period of time it will take to unwind the stock of imbalances, leading to a more prolonged period spent in the downturn. A more forward-looking policy rule places a higher penalty on a longer downturn, leading to stronger leaning at later stages of the financial cycle.

[Figure 7 here]

3.1.4 Persistence of the financial cycle

Consider next the optimal policy with a less persistent financial cycle process. Instead of a recurring Markovian boom-bust cycle characterisation, assume that financial imbalances tend to be self-correcting, following a stationary first-order autoregressive process:

$$f_t = \rho f_{t-1} + \gamma r_{t-1} + e_t^f \tag{11}$$

where $\gamma < 0$ and $\rho < 1$. The absence of a drift term and $\rho < 1$ imply that financial imbalances do not build up momentum in expectation. Rather they are mean-reverting and are expected to disappear over time even without a policy action. Any boom is only possible when there is a consecutive sequence of positive shocks pushing up f, an event that a priori has a very low probability (detailed results of this exercise are available in the online appendix).

These findings confirm the intuition that conventional AR parameterisations of the financial cycle (which imply much less persistence than seen in the historical booms and busts) imply much weaker incentives for early leaning than those associated with regime-switching persistence. Indeed, the expected benefit of pre-empting an unlikely tail event is small even when the AR parameter is nearly 1.

3.1.5 Endogenous downturn costs

One insight from Svensson (2017) is that the potential benefit of leaning can be fully offset if the cost of a crisis depends on the state of the economy. Intuitively, a weak economy is likely to suffer more than a strong economy for a given downturn. One way to incorporate this idea into the modelling framework is to alter the period loss function in the baseline in the following way

$$L_t = \left(y_t - cI(s_t = 2)\right)^2 \tag{12}$$

Under this specification, c can be interpreted as the output loss associated with a crisis. With this loss function, Svensson (2017) shows that leaning with the wind is the optimal policy. By leaning with the wind, a central bank can bolster the economy when a crisis is imminent so as to cushion the negative impact of a bust of a given cost.

The period loss function (12), however, has less appeal in our multi-period setting. Rewriting the baseline model with this loss function yields a fundamentally different optimal policy. In this case, the central bank can essentially offset the total impact of the bust on the loss function without a macroeconomic trade-off. In this case, the central bank lowers the policy rate until the output gap y is equal to c. Therefore, *the period loss is exactly equal to zero in the bust, ie cleaning up is costless*. With zero loss once the bust occurs, there is no incentive to lean against the wind during the boom. Figure 8 shows the optimal policy when using the period loss function (12), which entails substantial leaning *with* the wind as well as cleaning up during a bust.

[Figure 8 here]

A cleaning-up-after strategy is a first best strategy if it is feasible. This is the premise behind the benign neglect view – that the unwinding of financial imbalances can be costless in a macroeconomic sense. But the feasibility of such a strategy has been challenged by financial boom-bust experiences of the type associated with the GFC and others earlier in history.

An endogenous downturn cost is nonetheless an interesting issue. One can consider alternative loss functions that allow for costlier crises when the economy is weaker without necessarily implying a benign neglect strategy. For example, consider

$$L = -\frac{1}{2}y_t^2 - cI(s_t = 2)(1 + g(y_t)) \qquad (13)$$

where $g(y_t)$ is a kinked function, equal to zero for $y_t \geq 0$ and downward-sloping for $y_t < 0$. This loss function still implies that the downturn cost is higher the more negative is the output gap, but cannot be absorbed by a positive output gap. In other words, the downturn cost has a lower bound of c; namely there is an unavoidable cost that must be incurred once a crisis occurs.

The optimal policy associated with this case is shown in Figure 8, denoted by 'Cost with lower bound'. The function $g(y)$ is assumed to be linear for $y < 0$. The early leaning result is again established, and is qualitatively similar to the baseline case. The lesson we draw from this exercise is that as long as a central bank cannot completely eliminate the cost of a bust and the financial cycle is sufficiently persistent, then the optimal policy is to lean and to do so early in the cycle.

3.2 Adding a Phillips curve

A potentially important consideration is that of inflation. This section introduces inflation, draws a sharper distinction between real and nominal interest rates and examines the implications for our analysis. We consider two alternative assumptions about how real and nominal rates matter.

The first exercise assumes that the real interest rate gap influences both the (real) financial imbalances and output. The original model is augmented with an inflation gap equation in the macro block (ie a Phillips curve) and the inflation gap in the period loss function:

$$\pi_t = \beta_1^\pi \pi_{t-1} + \beta_2^\pi y_{t-1} + e_t^\pi \tag{14}$$

$$L_t = -\frac{1}{2}(y_t^2 + \pi_t^2) - cI(s_t = 2) \tag{15}$$

The output gap and financial cycle processes are as before, where f_t and r_t represent (real) financial imbalances and real policy rate gap, respectively. In this extended model, the key policy trade-off remains essentially the same. Leaning policy introduces a short-term cost to the policy maker, now including that of an inflation gap, but confers the benefit in terms of stabilising the financial cycle.

The optimal policy under this extended model is computed using the same set of calibrated parameters as in the baseline in order to focus on the marginal effects of introducing inflation. The Phillips curve parameters are calibrated to be consistent with the US inflation data ($\beta_1^\pi = 0.9$, $\beta_2^\pi = 0.037$, $\sigma_\pi = 0.33$). Figure 9 depicts the optimal policy in this extended model ('Phillips curve'), together with the previous baseline result ('Baseline'). As shown on the left panel, the optimal policy response to financial imbalances during the boom remains similar to before. There is a small reduction in the degree of leaning, due to the inflation cost of leaning. Figure 9 also shows another case where the Phillips curve is assumed to be more responsive to the output gap ('Steep Phillips curve', with $\beta_2^\pi = 0.1$). There is less leaning, which is to be expected. Qualitatively, the gist of previous conclusions carry over, including leaning early in the cycle.[10] The right panel shows the optimal policy response to inflation, which rises linearly with π_t (as expected given the quadratic loss function).

[Figure 9 here]

It might be argued that nominal and real interest rates could have different effects on the financial and real sectors. For example, the financial cycle process may be influenced by a 'search for yield' behaviour, reacting to the nominal interest rate rather than the real interest rate. An online appendix details this extension, showing that in this case the central bank may wish to lower the policy rate and 'lean with' the wind early in the cycle. The reason for this action is that leaning becomes more expensive under this search-for-yield assumption. For a given change in the real rate, the nominal rate tends to change by less because inflation and the real rate are negatively correlated in the model. For example, a tightening in the policy rate in the first stance leads to a lower output gap which then results in a decline in inflation. Lower inflation for a given real rate means that the nominal interest rate does not rise as much as the real rate and this, in turn, implies less restraint on the financial cycle. Put another way, to achieve the same level of restraint on the financial cycle, leaning would entail a higher macroeconomic cost under the 'search for yield' assumption than under the baseline case. However it is still the case that as the financial cycle matures, the central bank optimally leans against the wind.

[10] It may be argued that the definition of the crisis cost c should be broadened to include the decline in inflation associated with crises. This is not done here and, in this sense, provides a conservative estimate of the leaning benefit.

As a final thought on the role of inflation, it is instructive to compare our model to that of Svensson (2017) with his assumption that monetary policy has no long-run effect on the level of real debt (a form of 'money neutrality'). Under this assumption, Svensson (2017) argues that leaning only temporarily lowers debt in the short run at the expense of increasing it later on, an assumption that weakens the benefit of leaning. In our setting, while f_t does not converge to a fixed steady state, we maintain the same money neutrality assumption in that a temporary leaning policy does not have a permanent effect on financial imbalances. With or without leaning policy, f_t converges in distribution in the long run. With a systematic leaning policy, the long-run distribution of f_t is consistent with policy adjustments around the natural rate, implicit in our cyclical analysis. Indeed, both f_t and r_t may be interpreted as gap measures. The rationale for leaning in our analysis therefore does not rest on a deviation from money neutrality.

3.3 Assessing the bias with the marginal cost-benefit approach

As noted in the introduction, Svensson (2017) and IMF (2015) report evidence against the case for leaning using the marginal cost-benefit approach. In other words, they find that the marginal cost of a one-time tightening of the policy rate to rein in a financial boom exceeds the marginal benefits. It appears that their conclusion contradicts the findings in our model. This raises the question of whether the difference arises solely from some calibration detail or from the different approaches.

The latter could indicate that the full benefits of systematic leaning (that stabilises the endogenous financial cycle) are not captured in the marginal cost-benefit approach.

To address this possibility, it is instructive to first consider how the cost-benefit calculations differ between our approach and theirs. The optimal policy in our model is characterised by the first-order condition of Equation (7)

$$\begin{aligned}
&\frac{\partial}{\partial r_t}\left(\frac{1}{2}y_t^2\right) = \frac{\partial}{\partial r_t}E_t V(f_{t+1}, y_{t+1}, s_{t+1}) \\
&= \frac{\partial}{\partial r_t}\begin{pmatrix} P_t(crisis|boom)\int V\big(\tilde{f},\tilde{y},crisis\big)dF_t\big(\tilde{f},\tilde{y}\big) \\ +(1-P_t(crisis|boom))\int V\left(\tilde{f},\tilde{y},boom\right)dF_t\big(\tilde{f},\tilde{y}\big)\end{pmatrix}
\end{aligned} \tag{16}$$

which must hold *for all t*. The left-hand side is the marginal cost of leaning and the right-hand side is the marginal benefit of leaning.[11]

The characterisation of the marginal cost is the same under the marginal cost-benefit approach. The difference is the way in which the marginal benefit is measured. In the marginal cost-benefit approach, it is defined as

$$\frac{\partial P_t(crisis|boom)}{\partial r_t}\left(\tilde{V}(crisis)-\tilde{V}(boom)\right) \tag{17}$$

where P_t is the transition probability function from a boom to a crisis at *a given time t*. Note that the value function in Equation (17), $\tilde{V}$, is an estimate of the welfare cost of a crisis, based on the assumption of a no-leaning baseline. In general, an estimate of $\tilde{V}$ necessarily differs from V in Equation (16), which is a function of the optimal leaning policy rule. The latter takes into account the full effect of a systematic leaning rule on

[11] A more general statement of the value function approach of this problem with which to compare the two approaches is available in the online appendix.

the value function. In the case of a very persistent, endogenous financial cycle, the difference between $\tilde{V}$ and V is likely to be substantial, making the marginal benefit in Equation (17) a poor guide for the optimal policy.

One way to get a sense of the quantitative bias of the marginal cost-benefit approach is to apply the marginal cost-benefit approach to our model (see an online appendix). This quantitative assessment suggests that only by taking the full endogenous financial cycle dynamics into account and by considering optimal policy rules can the full benefit of a leaning policy be accurately measured.[12] It also highlights the fact that even in a model where systematic leaning is optimal, it still may not be welfare-enhancing to act in a discretionary fashion late in a financial cycle boom.

Suppose that both the output gap and interest rate gap are initially zero and the degree of imbalances is assumed to be $f = 2$, so that the risk of entering a downturn is already relatively high. The central bank then raises the policy rate gap for 4 quarters by 20 basis points, before reverting back to the output-targeting policy rule thereafter. Figure B.4 shows the paths of interest rate and output gaps, as well as the implied marginal cost and benefit of such an action. The marginal cost reflects the impact on output. The marginal benefit is the decline in the expected cost of a crisis, which is a product of (i) how much the rate increase curbs the rise in f, (ii) how much the transition probability is lowered and (iii) the downturn cost c.

[12] For example, central banks that apply the marginal cost-benefit approach in assessing their own policy forecasting models may be at risk of concluding erroneously that there are no net benefits to leaning when in fact there may be. Central banks would have to check whether the full set of dynamics of the crisis are modelled in order to evaluate whether the equivalence between the marginal cost-benefit approach and the full value-function approach is satisfied.

[Figure 10 here]

This one-time 20-basis-points tightening results in the marginal cost that is substantially larger than the marginal benefit. This result based on applying the marginal cost-benefit approach to our model is therefore consistent with that in the recent literature: a one-time tightening of monetary policy is not welfare-enhancing from the marginal cost-benefit analysis perspective! The type of leaning in this experiment is a one-time, discretionary use of the policy rate in the midst of a financial boom. In contrast, the type of leaning emphasised earlier in the paper is a policy of systematic rule-based leaning over the whole financial cycle.13

The value function approach and the marginal cost-benefit approach can yield similar leaning results if the policy rate does not materially influence the future value function. This would occur under two different assumptions. First, in our model, if the monetary authority can completely clean a financial bust's effect on the macroeconomy, the future value function is not affected by monetary policy actions during the boom. Then the marginal cost-benefit approach would yield accurate calculations. Also, if the financial cycle were immaterial or completely exogenous, the future value function would not be affected by leaning and the marginal cost-benefit approach would be accurate. But such assumptions about the financial cycle are at

[13] In fact, using this marginal cost-benefit approach, there is a case for leaning only when the magnitude is no greater than 5 basis points. This result is striking considering that the rule-based optimal policy derived earlier prescribes leaning as much as 80 basis points for the same initial conditions (see Figure 4). The benefits from leaning therefore are measured to be much lower quantitatively based on the one-time marginal cost-benefit calculation than based on the value function approach.

odds with the assumption of the exercise in this paper, ie that the financial cycle is costly, endogenous and unavoidable.

4 Conclusion

This paper establishes the case for a systematic leaning-against-the-wind policy in an environment where there is a recurring financial cycle with costly crises. The policy's main benefit arises from a dampening of the whole financial cycle – in terms of the frequency and severity of financial downturns. By dampening the cycle, a central bank reduces overall volatility, even though from a short-run perspective there may be somewhat higher macroeconomic conditional volatility.

The conclusion from our model stands in contrast to those models which downplay the dampening effect of monetary policy by assuming financial cycles that have a strong inherent tendency to self-correct. From our model's perspective, the existing analytical approach tends to underestimate the benefits of leaning by focusing on: i) strongly self-correcting financial cycles, and ii) a one-off policy action to address an imminent financial crisis. As a consequence, the policy assessment of a one-time discretionary tightening during a financial boom only gives a lower bound of the potential benefits from leaning.

In many respects, the modelling debate today harkens back to the one during the high and volatile inflation period of the 1970s and early 1980s. Model-wise, the debate then centred on the policy advice from short-term Keynesian models and from policy rule-based dynamic models (Lucas (1981)). Which model provided reliable guidance for policymakers? By the end of the debate, few believed that short-term discretionary policy actions to counter a rise in inflation was the best approach to achieve lasting price stability. Indeed, stop-go monetary policies in the 1970s proved to be ineffective and destabilising. Rather, strong price stability-oriented monetary policy frameworks, that were loosely "rule"-based and transparent, offered a more fruitful approach to achieving the desired results. In the same way, if financial cycles are considered an inherent part of the fabric of financially liberalised systems – which seems to be the case – models based on systematic policy responses may offer more reliable guidance that those based on one-off discretionary actions. And, if so, the debate should focus on understanding and assessing different policy rules. Of course, operationalising the basic thrust of such models and demonstrating the accuracy of the policy advice remain a challenge.

There are a number of practical reasons for qualifying the strength of our modelling conclusion. Given the difficulty of accurately tracking the financial cycle in real time, the net benefit from systematic leaning will depend on a consideration of type 1 and type 2 errors – that is, the errors of acting as if financial imbalances were growing when in fact they were not, and of not acting when financial imbalances were in fact growing. If a central bank cannot accurately track the extent of financial imbalances, additional costs need to be factored into the assessment of the net benefit. The consequences of such imperfect information conditions for cost-benefit assessment is left for future research.

Further analysis is also called for as our understanding of costly financial cycles evolves. In this paper, the financial cycle has been treated as being captured sufficiently well by a single variable, "f". Recent research, however, has emphasised a more nuanced relationship between different aspects of the financial cycle such as leverage, stocks versus flows of debt, debt service burdens, cross-border gross financial flows and bubbly prices. The dynamic interaction of these variables with the strength of the financial cycle and the economy is important. In addition, even though the trade-offs with inflation dynamics are addressed, more elaborate Philips curve relationships would allow us to explore an additional set of relevant policy issues. For example, uncertainties about the slope of the Philips curve can be analysed. As well, one could explore how policy credibility would be maintained even if leaning implied allowing inflation to run below target for a prolonged period of time. One possibility is to extend the notion of policy credibility beyond a fixed numerical target, and towards a more general through-the-cycle sustainability criterion.

In addition, an exploration of the micro mechanisms that generate endogenous financial cycles could yield important insights into the incentives that different actors in the financial system face and, as a consequence, the effectiveness of policy rates when leaning. Incorporating such features into our framework would enrich the analysis and shed new light on the nature of optimal policy.

Finally, even though we considered different ways of reflecting the costs of financial cycles, other important considerations may have been missing from our model. During the boom phase of a financial cycle, for example, there may be a rapid rise in the capital stock. On the one hand, as shown in recent studies for advanced economies (eg Gopinath et al (2017) and McGowan et al (2016)), the increase in the capital stock associated with booms has resulted in unproductive capital overhangs or inefficient capital allocation across economic sectors. By contrast, financial booms in our model do not have long-run negative consequences for productivity and potential growth. Such considerations could raise the benefits of systematic leaning (Juselius et al (2017) and Gourio et al (2018)).

On the other hand, financial booms may, under certain circumstances, contribute to future growth prospects. This may be the case in some emerging market economies. By accelerating the transition from a low to a high capital economy, a financial boom may yield additional benefits even if an eventual financial bust reverses some of them. Building in different types of growth effects may clarify some of the policy trade-offs facing central banks. In addition, the degree of substitutability or complementarity of macroprudential and monetary policy tools in managing the financial cycle is also key to any full calibration of the cost-benefit analysis. In this paper, we implicitly assumed that macroprudential tools could not fully stabilise the financial cycle (BIS (2016)).

Overall, our dynamic modelling approach is certainly not the last word on the issue of leaning against the wind. But the approach and the calibration of the financial cycle shows that our class of models provides a rich environment in which to assess the various costs and benefits associated with leaning.

References

Adrian, T and N Liang (2018): "Monetary policy, financial conditions, and financial stability", *International Journal of Central Banking,* vol 14(1).

Adrian, T and H S Shin (2010): "Liquidity and leverage", *Journal of Financial Intermediation*, vol 19(3).

Ajello, A, T Laubach, D Lopez-Salido, and T Nakata (2016): "Financial stability and optimal interest-rate policy," *Board of Governors of the Federal Reserve System Finance and Economics Discussion Series*, no 2016–067.

Alpanda, S and A Ueberfeldt (2016): "Should monetary policy lean against housing market booms?", *Bank of Canada Staff Working Paper*, no 2016–19.

Bank for International Settlements (2016): *86th Annual Report*, June.

Bernanke, B and M Gertler (1999): "Monetary policy and asset price volatility", in *Proceedings of the Federal Reserve Bank of Kansas City Economic Policy Symposium*.

——— (2001): "Should central banks respond to movements in asset prices?", *American Economic Review*, vol 91(2).

Blanchard, O (2000): "Bubbles, liquidity traps and monetary policy", in A Posen and R Mikitani (eds), *Japan's Financial Crisis*, International Institute of Economics.

Borio, C and M Drehmann (2009): "Assessing the risk of banking crises – revisited", *BIS Quarterly Review*, March.

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

Borio, C and H Zhu (2008): "Capital regulation, risk-taking and monetary policy: a missing link in the transmission mechanism?", *BIS Working Papers*, no 268.

Brunnermeier, M and M Oehmke (2012): "Bubbles, financial crises, and systemic risk", in G Constantinides, M Harris and R Stulz (eds), *Handbook of the Economics of Finance*, vol 2, Elsevier Science.

Cecchetti, S, H Genberg, J Lipsky and S Wadhwani (2000): *Asset prices and monetary policy*, International Centre for Monetary and Banking Studies.

Claessens, S, A Kose and M Terrones (2012): "How do business and financial cycles interact?", *Journal of International Economics*, vol 87(1).

Dell'Ariccia, G, L Laeven and R Marguez (2014): "Real interest rates, leverage, and bank risk taking", *Journal of Economic Theory*, vol 149(C).

Drehmann, M, C Borio and K Tsatsaronis (2012): "Characterising the financial cycle: don't lose sight of the medium term!", *BIS Working Papers*, no 380.

Filardo, A (1994): "Business-cycle phases and their transitional dynamics", *Journal of Business & Economic Statistics*, vol 12(3).

Filardo, A (2001): "Should monetary policy respond to asset price bubbles? Some experimental results", in G Kaufman (ed), *Asset Price Bubbles: Implications for Monetary and Regulatory Policies*, Elsevier Science.

——— (2004): "Monetary policy and asset price bubbles: calibrating the monetary policy trade-offs", *BIS Working Papers*, no 155. Published in Bank of Korea *Annual Conference Proceedings*.

——— (2006): "Asset Price Bubbles and Monetary Policy: A Multivariate Extension", unpublished BIS working paper, November.

——— (2011): "The Impact of the International Financial Crisis on Asia and the Pacific: Highlighting Monetary Policy Challenges from a Negative Asset Price Bubble Perspective", *BIS Working Papers*, no 356.

Gerdrup, K, F Hansen, T Krogh, and J Maih (2017): "Leaning against the wind when credit bites back", *International Journal of Central Banking*, vol 13(3).

Gopinath, G, S Kalemli-Ozcan, L Karabarbounis and C Villegas-Sanchez (2017): "Capital allocation and productivity in South Europe", *Quarterly Journal of Economics*, 132(4).

Gourio, F, A K Kashyap, and J Sim (2018): "The trade offs in leaning against the wind", *IMF Economic Review*, 66(1).

Gruen, D, M Plumb and A Stone (2003): "How should monetary policy respond to asset-price bubbles?", in *Asset prices and monetary policy*, a conference sponsored by the Reserve Bank of Australia, August.

Hamilton, J (2008): "Regime switching models", in *The New Palgrave Dictionary of Economics*, S Durlauf and L Blume (eds), Palgrave Macmillan.

International Monetary Fund (2015): "Monetary policy and financial stability", *IMF Staff Report*, September.

Jimenez, G, S Ongena, J Peydro and J Saurina (2012): "Credit supply and monetary policy: identifying the bank balance-sheet channel with loan applications", *American Economic Review*, vol 102(5).

Jorda, O, M Schularick and A Taylor (2013): "When credit bites back", *Journal of Money, Credit and Banking*, vol 45(2).

Judd, K (1999): *Numerical methods in Economics*, The MIT Press.

Juselius M, C Borio, P Disyatat and M Drehmann (2017): "Monetary policy, the financial cycle and ultra-low interest rates", *International Journal of Central Banking*, vol 13(3).

Laseen, S, A Pescatori and J Turunen (2015): "Systemic risk; a new trade-off for monetary policy?", *IMF Working Papers*, no WP/15/142.

Lucas, R (1981): "Econometric policy evaluation: a critique", in *Studies in Business-Cycle Theory*, MIT Press.

McGowan, M, D Andrews and V Millot (2016): "The walking dead? Zombie firms and productivity performance in OECD countries", *OECD Economics Department Working Papers*, forthcoming.

Miranda, M and P Fackler (2002): *Applied computational economics and finance*, The MIT Press.

Moreira, A and A Savov (2017): "The macroeconomics of shadow banking", *The Journal of Finance*, vol 72(6).

Rajan, R (2005): "Has financial development made the world riskier?", in *Proceedings of the Federal Reserve Bank of Kansas City Economic Policy Symposium*.

Riksbank (2013): "Financial imbalances in the monetary policy assessment", box in *Sveriges Riksbank Monetary Policy Report*, July.

Rudebusch, G and L Svensson (1999): "Policy rules for inflation targeting", in *Monetary Policy Rules*, J Taylor (ed), The University of Chicago Press.

Schularick, M and A Taylor (2012): "Credit booms gone bust: monetary policy, leverage cycles and financial crises, 1870–2008", *American Economic Review*, vol 102(2).

Svensson, L (2017): "Cost-benefit analysis of leaning against the wind", *Journal of Monetary Economics*, vol 90(C).

Vestin, D (2015): "Financial imbalances in the monetary policy decision", Sveriges Riksbank mimeo.

Woodford, M (2012): "Inflation targeting and financial stability," *Sveriges Riksbank Economic Review*, vol 1.

United States financial cycle

Figure 1

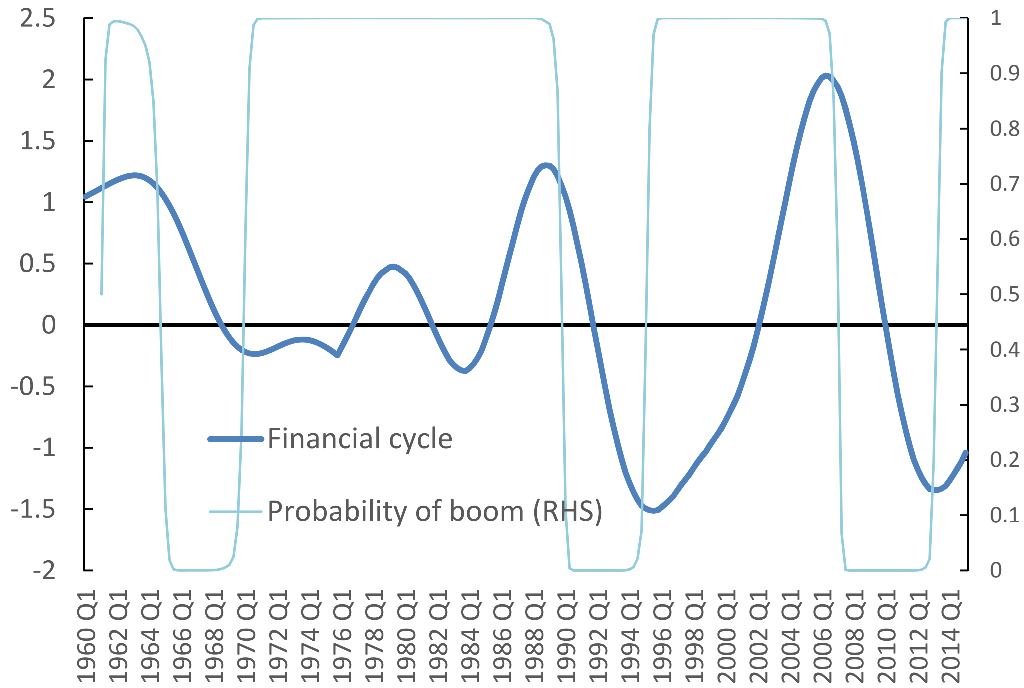

Endogenous financial cycle model estimates[1]

Table 1

	Scheme A	Scheme B	Scheme C
σ_f	0.0499*	0.0499*	0.0499*
	(0.000)	(0.000)	(0.000)
α_1	0.0422*	0.0424*	0.0425*
	(0.000)	(0.000)	(0.000)
γ_1	–0.0087*	–0.0087*	–0.0087*
	(0.000)	(0.000)	(0.000)
α_2	–0.1085*	–0.1082*	–0.1085*
	(0.000)	(0.000)	(0.000)
γ_2	0	0	0
ζ_1	0.2103*	0.1186*	0.1022*
	(0.000)	(0.000)	(0.000)
ψ_1	3.6182*	10.000*	10.000*
	(0.000)	(0.000)	(0.000)
ζ_2	0.2569*	0.1076*	1.0000*
	(0.000)	(0.000)	(0.000)
ψ_2	–2.0791*	–10.000*	–1.9675*
	(0.002)	(0.027)	(0.000)
f^H	1.5000*	1.1249*	1.1304*
	(0.000)	(0.002)	(0.004)
f^L	–1.1000*	–0.0140	–2.1480*
	(0.022)	(0.952)	(0.000)
Log likelihood	328.455	328.878	328.880

[1] Numbers in parentheses are the p-values, based on the likelihood ratio test that the parameters differ from zero. * indicates statistical significance at the 5% level.

Transition probabilities (estimated) Figure 2

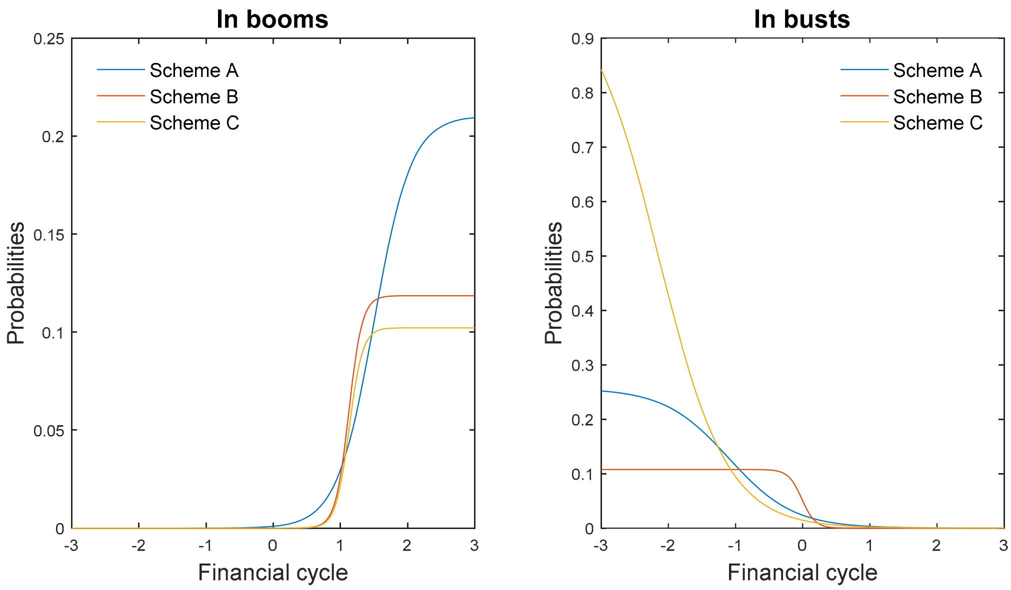

Optimal policy under Scheme A

Figure 3

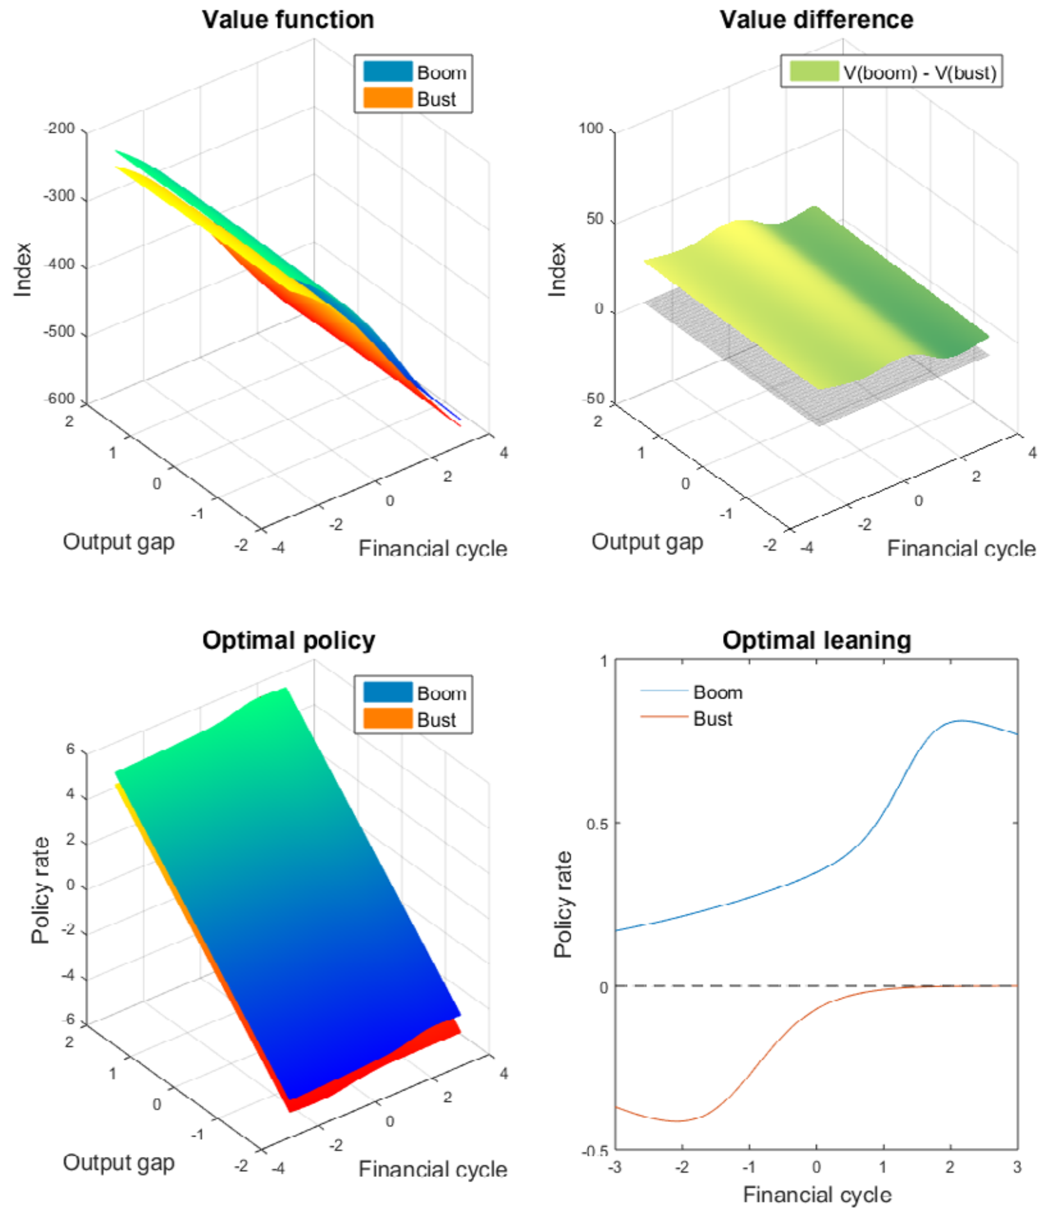

Optimal policy – A comparison across estimation schemes[1] Figure 4

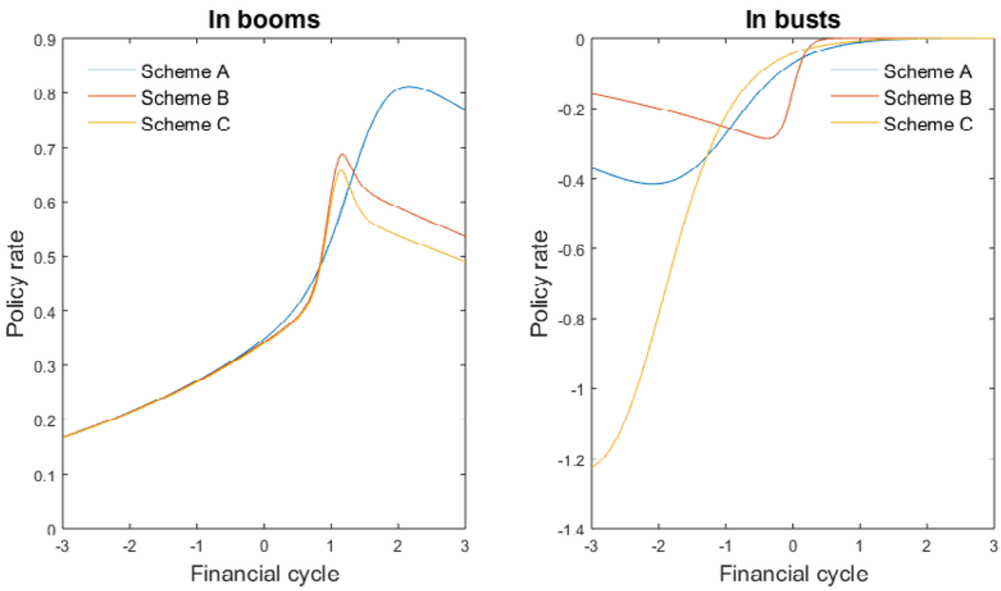

[1] The optimal policy is evaluated at $y=0$. It is also the "optimal leaning" for any level of output gap, namely the difference between the optimal policy and that which would have obtained if the central bank were to ignore the financial cycle dynamics.

Downturn costs and optimal policy[1] Figure 5

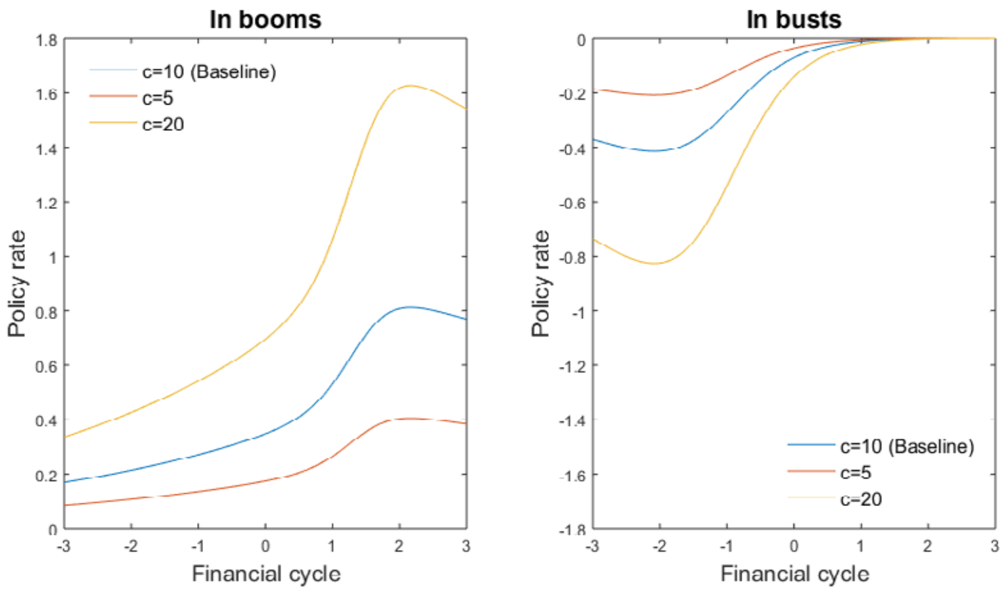

[1] Optimal policy evaluated at $y=0$, based on Scheme A.

IS curves and optimal policy[1]

Figure 6

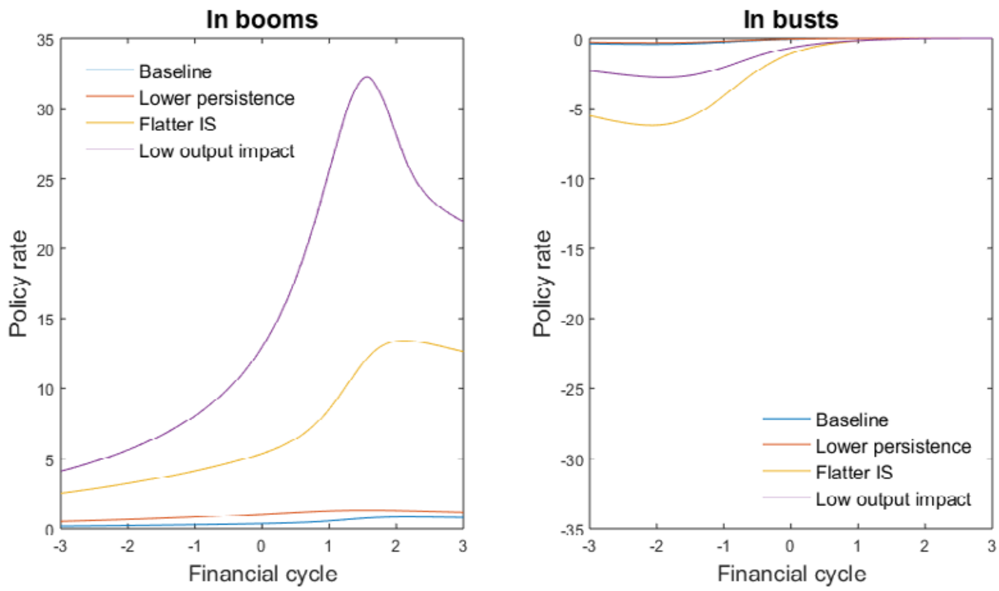

1 Optimal policy evaluated at $y=0$, based on Scheme A.

Discount factor and optimal policy[1] Figure 7

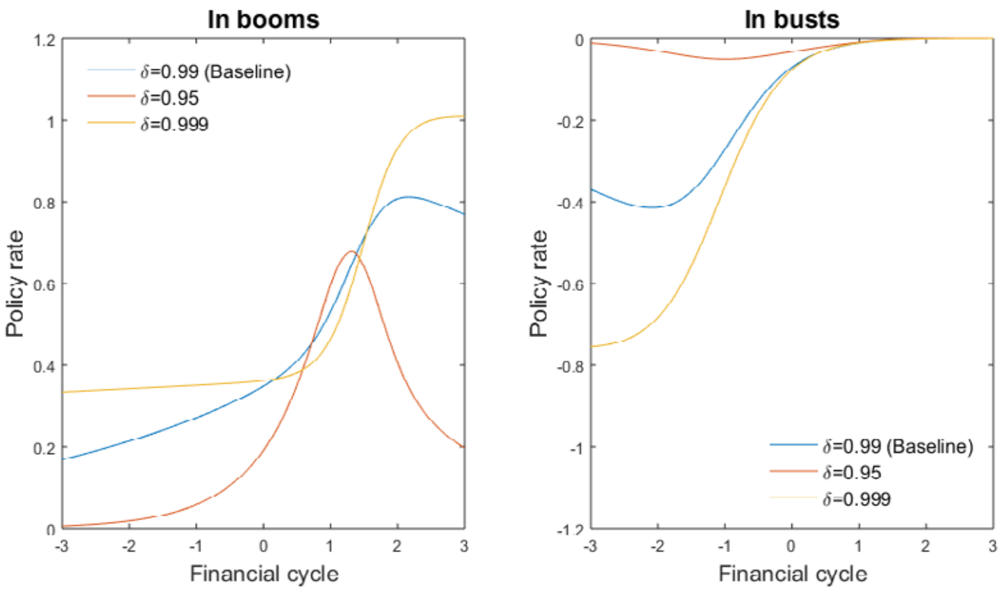

[1] Optimal policy evaluated at $y = 0$, based on Scheme A.

Downturn cost specifications and optimal policy[1]

Figure 8

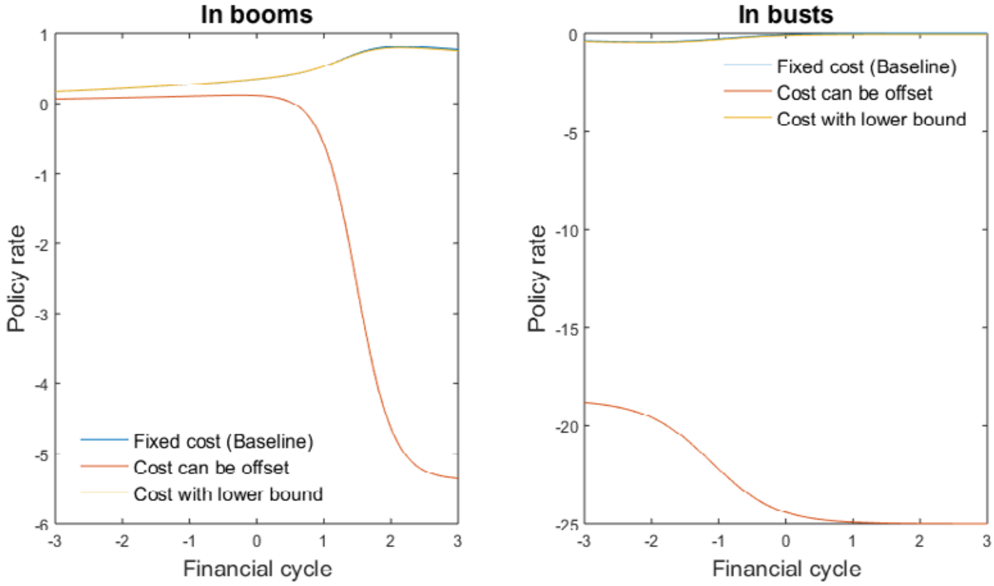

[1] Optimal policy evaluated at $y = 0$, based on Scheme A.

Adding a Phillips curve[1]

Figure 9

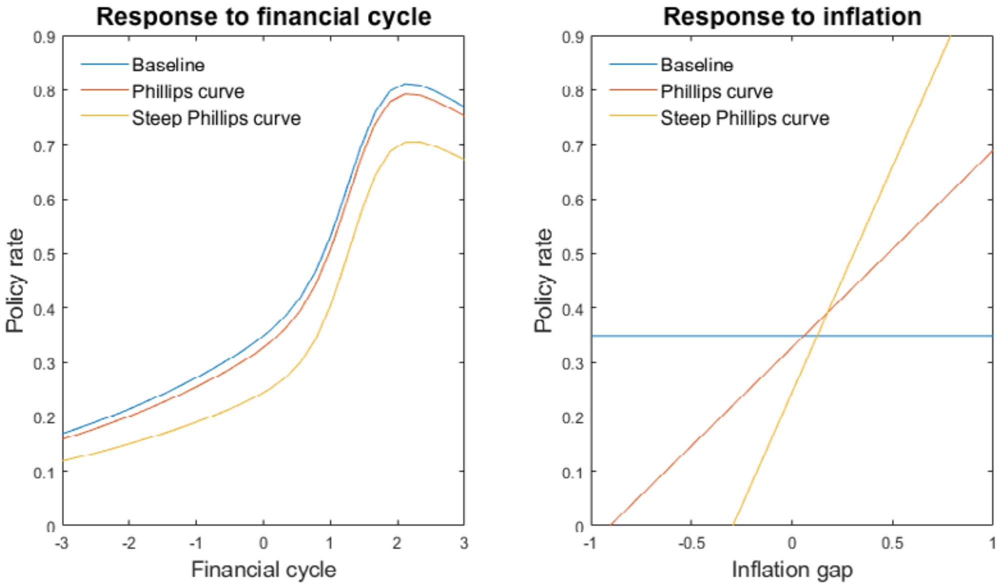

[1] Optimal policy during a boom phase. 'Baseline' represents optimal policy under Scheme A parameters. 'Phillips curve' represents optimal policy in the extended model with a Phillips curve, assuming the same parameters augmented by $\beta_1^\pi = 0.9$, $\beta_2^\pi = 0.037$, $\sigma_\pi = 0.33$. 'Steep Phillips curve' assumes $\beta_2^\pi = 0.1$. The left panel shows the optimal real rate r_t as a function of financial cycle f_t, evaluated at $y_t = \pi_t = 0$. The right panel shows the optimal real rate r_t as a function of inflation π_t, evaluated at $y_t = f_t = 0$.

Marginal cost-benefit analysis under the financial cycle model[1] Figure 10

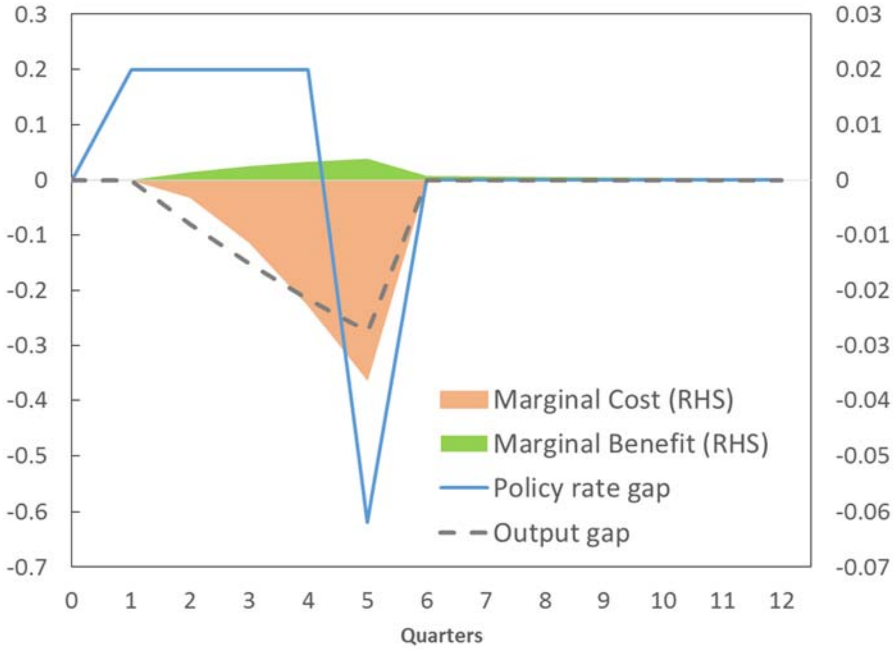

[1] The policy interest rate is assumed to be tightened by 20 basis points for 4 quarters, starting from the initial conditions of, $r = 0$, $y = 0$ and $f = 2$. The policy rate is then re-adjusted to target zero output gap subsequently. The marginal benefit is the change in the expected period downturn cost, based on estimation Scheme A. The marginal cost is the increase in quadratic loss as a result of lower output gap.

Online appendix A – Estimation procedure for the financial cycle block

Let Ω_{t-1} denote the information set at time $t-1$ and let $\theta = \{\sigma_f, \alpha_1, \alpha_2, \gamma_1, \gamma_2, \zeta_1, \zeta_2, \psi_1, \psi_2, f^H, f^L\}$ be the set of the financial cycle parameters. At time $t-1$, the likelihood function of observing f_t conditional on regime $s_t = s$ is given by

$$g(f_t|s, \Omega_{t-1}, \theta) = \frac{1}{\sqrt{2\pi\sigma_f^2}} \exp\left[-\frac{(f_t - f_{t-1} - \alpha_s - \gamma_s r_{t-1})^2}{2\sigma_f^2}\right]. \quad (1)$$

Denoting the conditional probability of being in regime s by $\xi_{s,t-1}$. The unconditional likelihood function is then given by

$$g(f_t|\Omega_{t-1}, \theta) = \sum_{\tilde{s}} \sum_{s} \xi_{s,t-1} p_{t-1}(\tilde{s}|s) g(f_t|\tilde{s}, \Omega_{t-1}, \theta). \quad (2)$$

To complete the algorithm and generate the likelihood for all t, the conditional regime probability is recursively updated according to

$$\xi_{\tilde{s},t} = \frac{\sum_s \xi_{s,t-1} p_{t-1}(\tilde{s}|s) g(f_t|\tilde{s}, \Omega_{t-1}, \theta)}{g(f_t|\Omega_{t-1}, \theta)}. \quad (3)$$

The maximum likelihood estimates are calculated numerically as the solution to

$$\hat{\theta}_{MLE} = argmax_\theta \left(\sum_{t=1}^{T} \log g(f_t|\Omega_{t-1}, \theta) \right), \quad (4)$$

subject to the restrictions $\alpha_1 \geq 0$, $\alpha_2 \leq 0$, $\gamma_1 \leq 0$, $\gamma_2 \leq 0$, $\psi_1 > 0$, $\psi_2 < 0$, $0 < \zeta_1 < 1$, $0 < \zeta_2 < 1$, $f^H > 0$, $f^L < 0$.

Online appendix B – Detailed robustness results

B.1 Value functions under estimation schemes B and C

As mentioned in the main text, the shapes of value functions and optimal policies are qualitatively the same across all estimation schemes. The value function corresponding to the boom strictly dominates that under the bust. For any regime, the value function is decreasing in the degree of financial imbalances. Optimal policy prescribes leaning during the boom.

Optimal policy under Scheme B

Figure B.1.1

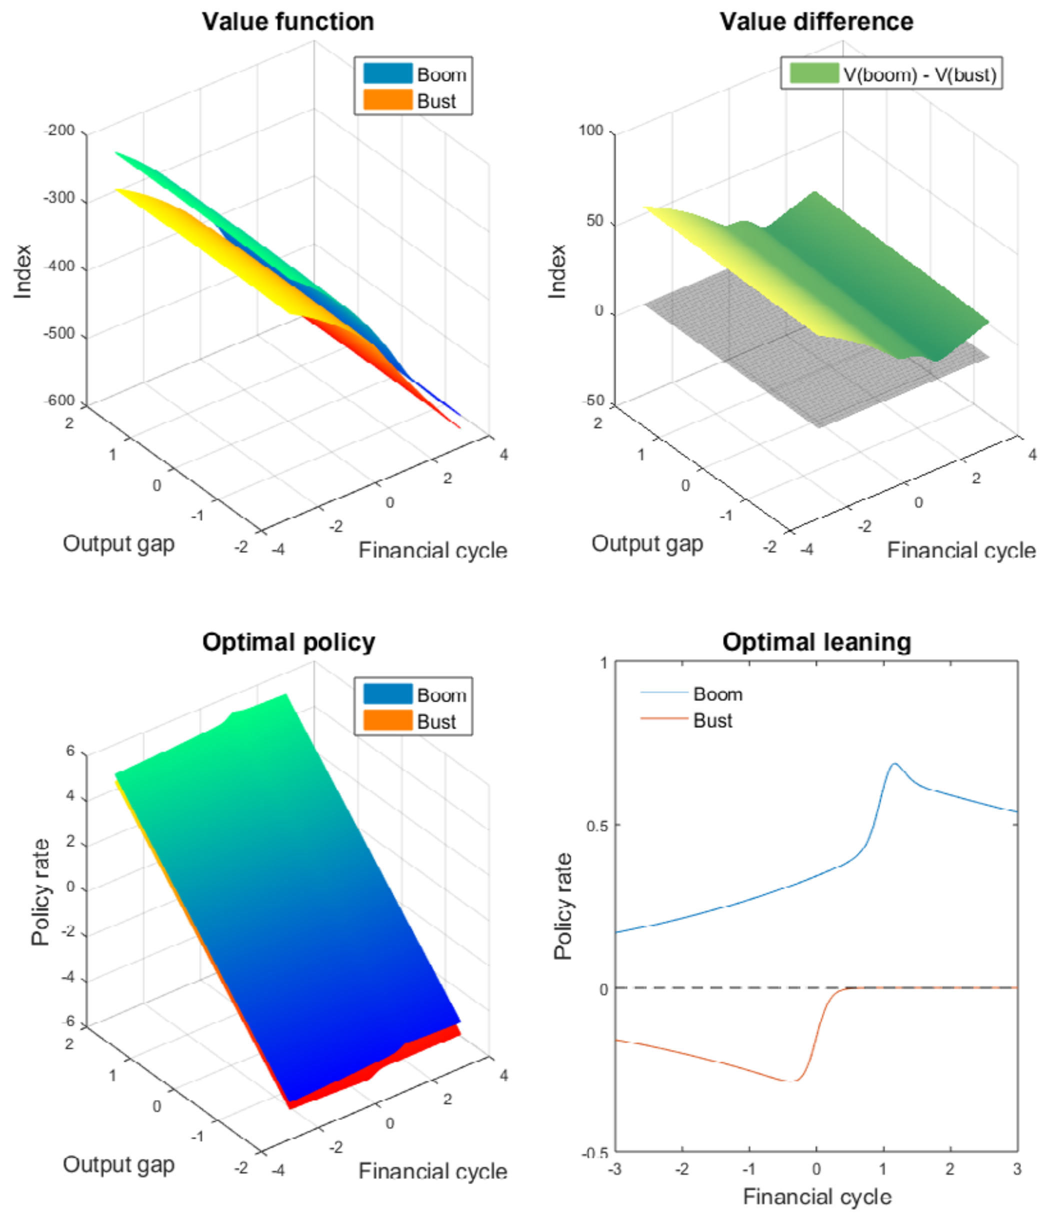

Optimal policy under Scheme C

Figure B.1.2

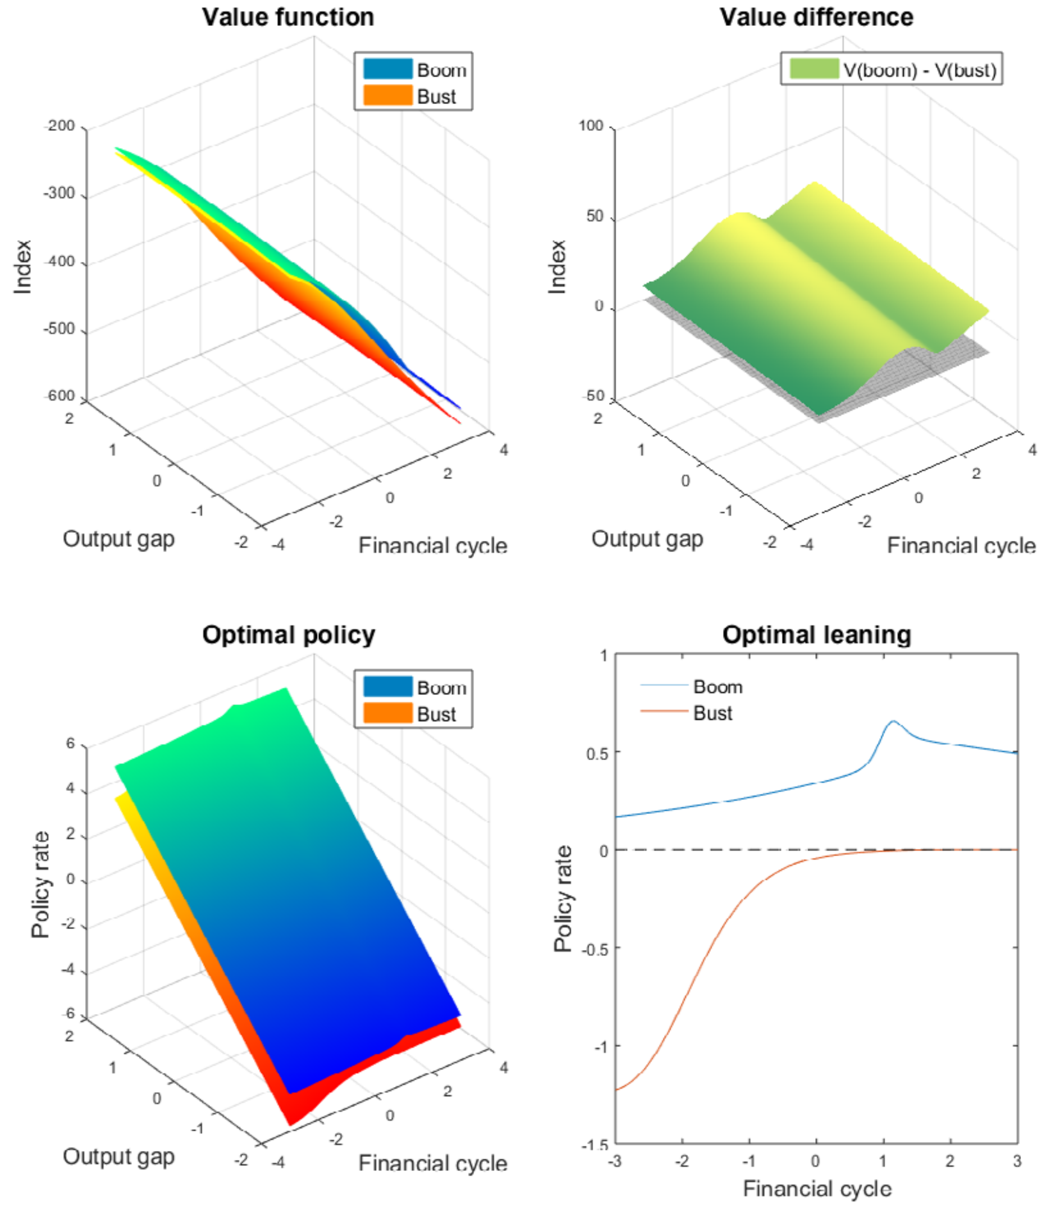

B.2 Alternative financial cycle processes

This robustness exercise assumes that the financial cycle process f follows a stationary first-order autoregressive process:

$$f_t = \rho f_{t-1} + \gamma r_{t-1} + e_t^f$$

Figure B.2.1 shows the optimal policies when $\rho = 0.9$ and $\rho = 0.99$. [1] In both cases, there is little incentive for early leaning, because the central bank faces a much less favourable trade-off between the short-term macroeconomic costs of leaning and the impact on the likelihood of a crisis. The more persistent process with $\rho = 0.99$ entails earlier leaning, as it takes longer for financial imbalances to unwind. A higher persistence implies that the likelihood of a future crisis is larger at any given f. However, when f becomes very high, the looming crisis provides an incentive for leaning more strongly in either cases.[2]

Financial cycle processes and optimal policy[1] Figure B.2.1

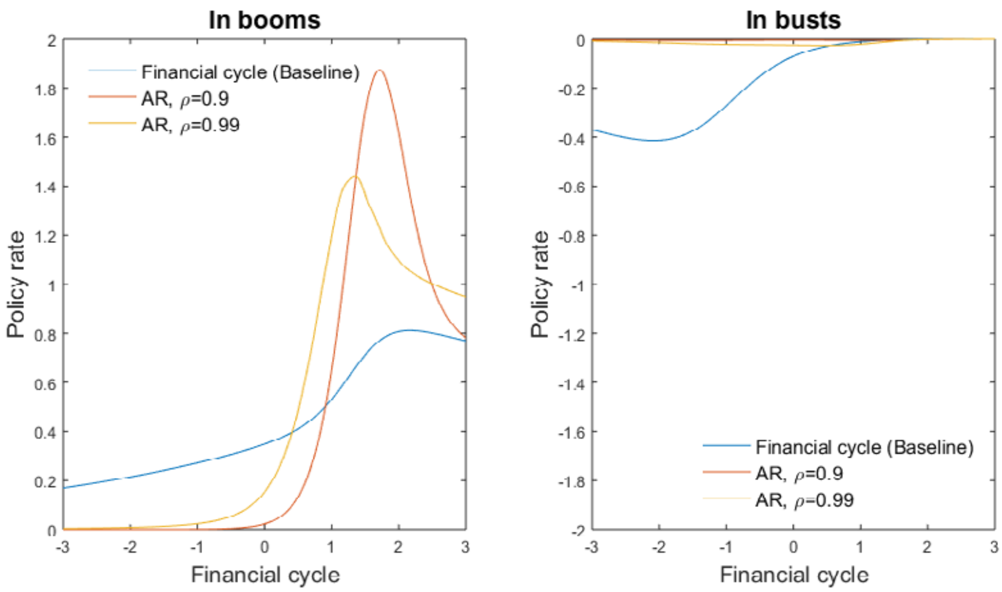

[1] Optimal policy evaluated at $y = 0$, based on Scheme A.

[1] As in the baseline, it is assumed that the probability of the boom going bust is increasing in f_t, described by a logistic function. Once a bust ensues, it is assumed as before that f_t must decline persistently until there is a regime switch.

[2] The case of vanishing persistence $\rho \to 0$ (not shown) implies less leaning than any of the cases presented. There is still a gain from leaning in such a case, even if the financial cycle variable is close to being an i.i.d variable. This is because the central bank wants to increase the probability of entering a bust with low imbalances, which would entail a less prolonged bust phase.

B.3 Introducing the Phillips curve with a nominal search-for-yield assumption

The financial cycle could be influenced more by the nominal rather than the real interest rate. This would be the case, for example, when the financial cycle is being driven by 'search for yield' behaviour. To consider the implications of this possibility, we assume that f_t depends on the nominal interest rate i_t, while the output gap and inflation continue to depend on the real interest rate $i_t - \pi_t$. The Phillips curve extension of the model can be modified in the following way:

$$f_t = f_{t-1} + \alpha_{s_t} + \gamma_{s_t} i_{t-1} + e_t^f \tag{5}$$

$$y_t = \beta y_{t-1} + \phi(i_{t-1} - \pi_{t-1}) + e_t^y \tag{6}$$

The optimal policy is reported in Figure B.3.1 (together with the two earlier results for both the baseline and the baseline augmented with a Phillips curve). The optimal policy in this case (labelled 'Search for yield') implies a lower interest rate at every stage of the financial cycle. The yellow line in the figure illustrates this point. The reason is that leaning becomes more expensive in the case of the search-for-yield assumption. For a given change in the real rate, the nominal rate tends to change by less because inflation and the real rate are negatively correlated in the model. For example, a tightening in the policy rate in the first stance leads to a lower output gap which then results in a decline in inflation. Lower inflation for a given real rate means that the nominal interest rate does not rise as much as the real rate and this, in turn, implies by Equation (5) less restraint on the financial cycle. Put another way, to achieve the same effect on the financial cycle, leaning would entail a higher macroeconomic cost under the 'search for yield' assumption than under the baseline case. Notice however that as the financial cycle matures, the central bank optimally leans against the wind.

The decline in the yellow line is sufficiently large to suggest leaning "with" the wind is optimal at early stages in the financial cycle. Why? There are two key reasons. First, consider an economy where the inflation and output gaps are zero. Given concerns about the financial cycle in the future, the central bank would prefer, all else the same, somewhat tighter financial conditions. Instead of raising the real policy rate, which as noted above is more expensive, the central bank has an incentive to boost inflation over time to raise the nominal interest rate as a means to sustainably tighten conditions on the financial cycle. To do this, the output gap must increase and that is achieved initially by a lowering the real policy rate to drive up inflation over time. Second, since boosting output and inflation helps cushion the economy against the macroeconomic costs of future leaning as the financial cycle matures, there is an additional incentive to implement a lower policy rate at an earlier stage of the financial cycle.

Phillips curve version of model with a nominal search-for-yield assumption[1] Figure B.3.1

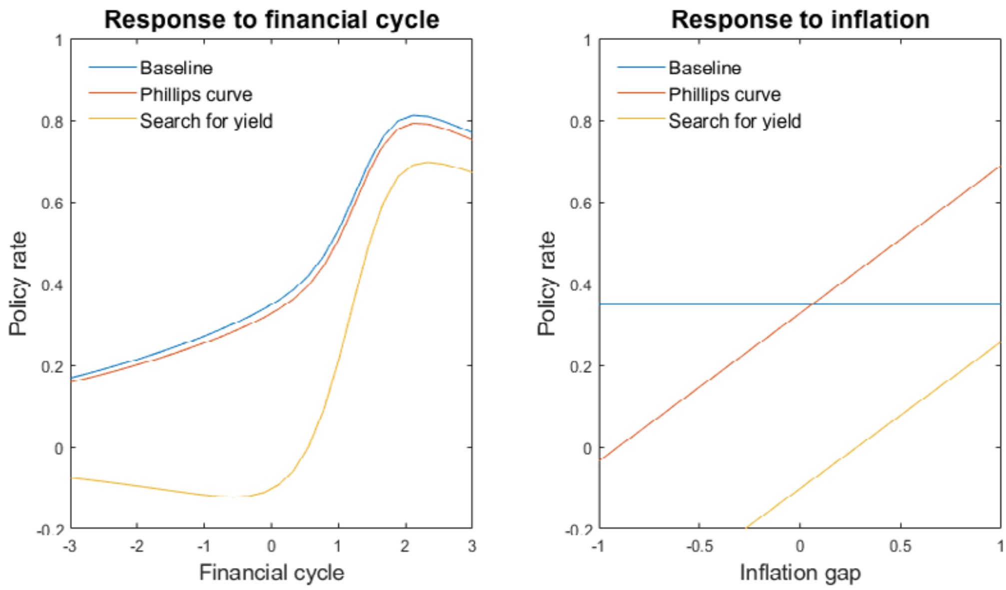

[1] Optimal policy during a boom phase. 'Baseline' represents optimal policy under Scheme A parameters. 'Phillips curve' represents optimal policy in the extended model with a Phillips curve, assuming the same parameters augmented by $\beta_1^\pi = 0.9$, $\beta_2^\pi = 0.037$, $\sigma_\pi = 0.33$. Under 'Search for yield', the financial cycle is assumed to depend on the nominal rather than real interest rate. The left panel shows the optimal real rate r_t as a function of financial cycle f_t, evaluated at $y_t = \pi_t = 0$. The right panel shows the optimal real rate r_t as a function of inflation π_t, evaluated at $y_t = f_t = 0$.

Online appendix C - General statement of the monetary policy problem

This appendix highlights key insights about rule-based leaning that can be gleaned from examining a general statement of the monetary policy problem addressed in the paper. The insights reinforce the gist of the robustness section that our conclusions are not a result of specific parameterisations. Rather, our conclusions reflect the dynamic, rule-based way in which we define the monetary policy problem being faced by central banks.

At the core of the problem, a central bank attempts to stabilise the macroeconomy by setting a time path for the policy interest rate. Conventionally, the policymaker's objective is to optimise a quadratic function of the period loss,

$$-\tfrac{1}{2} E \sum_{t=0}^{\infty} \delta^t \, [\pi_t^2 + \lambda y_t^2], \tag{7}$$

subject to the equations describing the macroeconomic block and the financial cycle; π_t and y_t denote output and inflation gaps, δ is the discount factor.

For pedagogical reasons, we re-cast this optimal control problem as a dynamic programme to highlight ways in which the financial cycle may matter for policy. It is instructive to compare the standard pre-crisis model to those which explicitly include a financial cycle.

Starting with a standard pre-crisis model, it focused on inflation and output with no role for the financial cycle. The problem can be written in the following way:

$$V(y, \pi) = \max_{r_1(y,\pi)} \left\{ -\frac{1}{2} (\pi_t^2 + \lambda y_t^2) + \delta E V(y', \pi') \right\}, \tag{8}$$

subject to the law of motion of the macroeconomy. In this formulation, the role of the financial cycle is absent reflecting a predominant pre-crisis view that finance could reasonably be treated as acting largely 'as a veil'. As well, the interest rate rule $r_1(y, \pi)$ is assumed to be of the Taylor-type. Solutions to this formulation are well known in the literature (see eg Woodford (2012)).

The Great Financial Crisis raised awareness of financial sector developments which required new modelling efforts. We consider two possible ways in this appendix to enrich the analysis and raise the prominence of macro-financial interactions via the financial cycle : i) augment the model to include a financial cycle block; ii) include an explicit reaction of the central bank to the financial cycle.

The first way to address the financial cycle implies re-writing the value function as

$$V(y, \pi, f, s) = \max_{r_2(y,\pi)} \left\{ -\frac{1}{2} (\pi_t^2 + \lambda y_t^2) + \delta E V(y', \pi', f', s') \right\}, \tag{9}$$

where central bank chooses a policy rule as before, ie being a function of only output and inflation, but the financial cycle f is included explicitly in the laws of motion of the equation system. Note that the financial cycle specification can be flexible enough to include financial cycle state-dependent outcomes (captured by the crisis indicator $s \in \{boom, bust\}$). Equation (9) indicates that the central bank's policy rule, $r_2(y_t, \pi_t)$, only incorporates reactions to the financial cycle to the extent that it influences output and inflation.[3]

The second way to address the central bank modelling challenge is to consider policy rules which depend directly on the state of the financial cycle f and regime s. The monetary policy problem can be recast in the following way,

$$V(y, \pi, f, s) = \max_{r_3(y,\pi,f,s)} \left\{ -\frac{1}{2}(\pi_t^2 + \lambda y_t^2) + \delta EV(y', \pi', f', s') \right\}, \tag{10}$$

where the central bank responds directly to the financial cycle in the policy rule.

Comparing this approach to recent modelling efforts.

Svensson (2016) and the IMF (2015) explore the trade-offs associated with leaning using an approach for calibrating the costs and benefits of leaning. The benefits of their approach include the ease of computation and innovative modelling of financial stability risks. Leaning in their approach reduces the probability of a crisis; but it comes at the short-term cost of lower economic activity and inflation. While insightful, the approach implies various restrictions when compared to Equation (10). Given these restrictions, the approach may underestimate the full benefits of leaning to the extent that the restrictions are not justified.

This *discretionary marginal cost-benefit approach* can be seen as imposing the following restrictions on the value function,

$$\begin{aligned} EV(y', \pi', f', s') &= \int V^1(y', \pi') g(y', \pi') V^2(f', s') h(f', s') \\ &= \int \left\{ \begin{matrix} p(crisis, f') V^1(y', \pi') g(y', \pi') V^2(f', crisis) \tilde{h}(f') \\ +(1 - p(crisis, f') V^1(y', \pi') g(y', \pi') V^2(f', boom) \tilde{h}(f') \end{matrix} \right\} \end{aligned} \tag{11}$$

Recasting the policy problem in this way highlights the impact of the restrictions in three key ways.

First, the discretionary marginal cost-benefit approach focuses on one upcoming potential crisis. This is equivalent to assuming that once a crisis occurs, the economy is crisis-free forever afterwards. This imposes the restriction that the probability of an imminent crisis, $p(crisis, f')$, becomes zero after the realisation of a crisis and $1 - p(crisis, f')$ becomes 1. In other words, the central bank is assumed to ignore the possible far-reaching impact of different policy rules on the future evolution of the financial cycle; this is equivalent to dropping $V^2(f', s')h(f', s')$ out of the Equation (11). Such restrictions may be reasonable if a crisis is truly a once in a lifetime event. But this

[3] Note, in particular, that the monetary policy reaction function in this case is restricted to be of the standard Taylor-type; it only responds directly to output and inflation, but not directly to the financial cycle itself. This does not mean the policy implications are similar to the pre-crisis setup. Quite the contrary, the optimal policy will differ, owing to the influence of the financial cycle on output and inflation.

may not be reasonable if crises are endogenous recurring events as modelled in the paper.

Second, the discretionary marginal cost-benefit approach may not fully capture the notion that systematic policy rules can fundamentally reshape the endogenous financial cycle. Different rules can affect the evolution of the (reduced-form) probabilities and hence the time series behaviour of the financial cycle.[4] Specifically, the discretionary marginal cost-benefit approach implies that the shape of the expected value function $EV(y, \pi, f, s)$ is invariant to the impact of different policy interventions. This criticism is a form of the classic Lucas critique and hence raises questions about measuring the full range of benefits from a rule-based leaning monetary policy.

Finally, the discretionary marginal cost-benefit approach may also underestimate the benefits of leaning when the cost of crises vary over time. Consider the case when the cost of a crisis is assumed to be fixed over time (ie under the strong assumption that the cost is not directly affected by the state of the macroeconomy) and when the cost only shows up in the loss function of the central bank. This implies that the future value function $EV(y, \pi, f, s)$ is separable into the sub-value functions $V^1(y', \pi')g(y', \pi')$ and $V^2(f', s')h(f', s')$. The first term is the sub-value function and captures the impact of the macroeconomic block on the value function, and the second term captures the financial cycle cost on the value function.[5] In more general cases, there would be interactions between the macro variables and an endogenous financial cycle that would preclude such separability, opening up the possibility of even further gains from leaning.[6]

All these implied restrictions underscore the fact that the discretionary marginal cost-benefit approach may not capture the full benefits of leaning. It also suggests that *systematic* leaning may outperform discretionary, one-time policy actions. Ultimately, the appropriateness of each approach is an empirical issue. We have shown in the paper that the value function approach with an endogenously determined financial cycle that is calibrated to the historical data yields evidence in favour of systematic leaning.

---

4 In the general statement of the problem, a shift from a non-leaning policy rule to a leaning policy rule can alter the shape of $EV(y', \pi', f', s)$ and hence the likelihood of future crises. As a result, different policy rules have different implications for the (reduced form) amplitude and duration of the financial cycle.

5 It is useful to note that if $V^2(f', s')h(f', s')$ is a constant, eg the policy rate does not affect the financial cycle, then there would be no reason to lean against the wind. Another special case of this assumption is benign neglect; if the central bank can costlessly clean up a crisis, then the optimisation problem simplifies to a 'finance as a veil' monetary policy problem (ie without a financial cycle).

6 Some versions of the discretionary marginal cost-benefit approach assume that the crisis cost is fixed in size. It is sometimes suggested that the fixed cost could reflect the present value of all the costs associated with a crisis. However, such restrictions rule out the state-dependence possibility that longer and more pronounced financial cycles lead to larger and more costly financial crises. The benefits of leaning with state-dependent costs could result in larger benefits to leaning.