Korsgaard, Søren

Working Paper

Ahead of the cycle

Danmarks Nationalbank Working Papers, No. 143

Provided in Cooperation with:
Danmarks Nationalbank, Copenhagen

Suggested Citation: Korsgaard, Søren (2019) : Ahead of the cycle, Danmarks Nationalbank Working Papers, No. 143, Danmarks Nationalbank, Copenhagen

This Version is available at:
https://hdl.handle.net/10419/227861

WORKING PAPER

DANMARKS NATIONALBANK

17 OCTOBER 2019 — NO. 143

Ahead of the Cycle

Søren Korsgaard
sk@nationalbanken.dk
DANMARKS NATIONALBANK

Ahead of the Cycle

Abstract
I examine the conduct of macroprudential policy in an environment where economic indicators move in a cyclical fashion, policy works with a lag, and there are adjustment costs to changing policy. In this setting policy instruments such as the countercyclical capital buffer should be set not only based on the present state of the cycle, but also on where the cycle is expected to be in the future and on the current level of buffer.

Resume
Jeg analyserer, hvordan makroprudentiel politik bør implementeres, hvis økonomiske indikatorer er cykliske, makroprudentiel politik virker med forsinkelse, og der er tilpasningsomkostninger forbundet med at ændre politik. I denne situation bør politikinstrumenter som den kontracykliske kapitalbuffer blot være en funktion af den nuværende cykliske situation, men også af hvor cyklen er på vej hen og det gældende niveau for bufferen.

Key words
Macroprudential policy; financial stability

JEL classification
G28; E58

Acknowledgements
The authors wish to thank colleagues from Danmarks Nationalbank.

The authors alone are responsible for any remaining errors.

Ahead of the Cycle

Søren Korsgaard*, Danmarks Nationalbank

October, 2019

Abstract

I examine the conduct of macroprudential policy in an environment where economic indicators move in a cyclical fashion, policy works with a lag, and there are adjustment costs to changing policy. In this setting policy instruments, such as the countercyclical capital buffer, should be set not only based on the present state of the cycle, but also on where the cycle is expected to be heading in the future and on the current buffer level.

Keywords: Macroprudential policy; financial regulation; countercyclical capital buffer

JEL-classification: G28; E58

*E-mail: sk@nationalbanken.dk. I would like to thank colleagues at Danmarks Nationalbank for valuable feedback.

1. Introduction

Macroprudential authorities spend considerable efforts looking at economic and financial indicators, trying to identify the state of various cycles and deciding on the appropriate level of macroprudential policy instruments such as the countercyclical capital buffer. At the same time they must take into account the costs of changing policy. These costs can be thought of as adjustment costs on the part of the agents affected by the policy, e.g. banks having to raise capital or choosing to reduce assets, or simply as the costs of implementing the policy. In some jurisdictions, for example, one authority is responsible for making recommendations on macroprudential policy, but must persuade another authority to implement the policy (ESRB, 2018). In addition, policy makers must contend with the fact that their policies are sometimes implemented with a considerable lag, at which point the cycle might be in a different state from when the policy decision was made.

I construct a model which captures these basic features of macroprudential policy making in a tractable framework. It features a policy maker who, based on a range of indicators, has a view on the ideal, or "target", levels of various policy instruments. For concreteness, think of a single instrument such as the countercyclical buffer rate. If there were no frictions, the buffer rate would be set equal to the target. In the model there are frictions. It is costly to adjust policy, and the buffer rate set in the current period only takes effect in the next period. The target moves in a cyclical (and stochastic) fashion, and the policy maker trades off the cost of straying from the target against the cost of adjusting the buffer.

The key result of the paper is that policy makers should consider not just the state of the cycle today, but also where it is heading in the future, the current level of policy instruments, and how costly it is to change policy. If the model were static, the optimal rate would be a weighted average of the current buffer and the optimal (target) buffer. Because policy makers must take into account the movement of the target and their own future actions, however, they need to think ahead of the cycle. Specifically, I show that the optimal buffer can be expressed in closed form as a linear combination of the buffer in place, the existing target, and the expected value of the target in the next period.

The principles underlying the model are by no means original to macroprudential policy, but have a wide variety of applications in other areas ranging from the design of missile systems to sports where agents also "aim ahead of the target".[1]

There are two essential steps to solving the model. The first step is to show that the model can be cast in linear-quadratic form[2], which establishes the existence and uniqueness as well as the form of the solution. The second step is to separate the state variables into two groups, (1) existing buffer levels and (2) past and present levels of economic indicators, and then split the value function into two parts involving these terms, respectively, and a third term involving cross-terms. This allows one to write the problem in a manner which is mathematically equivalent to that in Gârleanu and Pedersen (2013) and one can then directly apply the closed-form solutions identified in that paper.

I also consider simple numerical extensions to the model, each intended to illustrate a friction in actual policy making. The first extension is to vary the duration of the implementation lag. A longer lag simply implies that policy makers must think even further ahead in the cycle. I then proceed to look at asymmetries in policy making. While it may be costly to increase capital levels for banks, it is not costly for banks to reduce capital levels. This asymmetry implies that while the build-up of capital buffers should be gradual, the release of capital buffers can be swifter. Another asymmetry is that while increases in instruments such as the countercyclical capital buffer only take effect after an extended period, e.g. 12 months, the instruments can be released with immediate effect. In addition, the cost of building capital is likely to be state-dependent and, specifically, cheaper in "good" states such as upturns. The effect of this asymmetry is to have a swifter build-up of capital buffers when the underlying cycle is improving.

While stylized and intended as a story-telling device, the model captures key elements of actual policy. The evolution of economic and financial indicators clearly affects policy decisions. A number of macroprudential authorities also emphasize that buffers should be

[1] In fact, the inspiration for the paper was Gârleanu and Pedersen (2013) who apply the same principles to trading strategies in financial markets.

[2] See chapter 5 in Ljungqvist and Sargent (2012) for a discussion of results related to linear-quadratic dynamic programming.

built-up both *ahead of the cycle* and *gradually*, exactly the "prescriptions" laid out in this paper. For example, in a policy statement on its strategy for setting the countercyclical capital buffer, the Bank of England (2016) lays out a "moving early" strategy and emphasizes how it intends to vary the buffer gradually to reduce its economic cost. One of the Bank of England's five core principles reads: "By moving early, before risks are elevated, the FPC expects to be able to vary the CCyB gradually, and to reduce its economic cost." Other macroprudential authorities have similar strategies which mirror the principles of the paper. In their memo on the countercyclical capital buffer, the Danish Systemic Risk Council (2017) emphasize an early and gradual build-up of the buffer.

In terms of related literature, there is extensive research on the rationales for macroprudential policy (see, amongst others, Farhi and Werning (2016)), its effects (e.g. Jiménez et al. (2017) and Galati and Moessner (2018)), and institutional details and applications (Claessens (2018) and Lim et al. (2011)). Some papers, for example Bianchi and Mendoza (2018), derive optimal macroprudential policies, but these are quite different from - and more complex than - the macroprudential policy instruments actually used. There is also a number of more practical papers, such as Drehmann et al. (2010) and Drehmann et al. (2011), which address which indicators to use when setting the levels of policy instruments such as the countercyclical capital buffer. The main contribution of this paper is to highlight that policy making should be forward-looking and not just concerned with the current values of such indicators, but also their projected future evolution.

The rest of the paper proceeds as follows. Section 2 outlines the model in the simple case where a policy maker chooses a single instrument (a "buffer rate") to trade off the costs of straying from a moving target and of adjusting policy. Section 3 describes the more general case with multiple policy instruments and targets, and the targets themselves reflect underlying financial or economic indicators. In section 4, I numerically examine the effects of introducing realistic frictions to the model. Section 5 concludes.

2. A simple model of macroprudential policy

I start out by considering a simple, motivating example in a static setting. Consider a policy maker who inherits a buffer rate b_{t-1} and must choose a new buffer rate b_t. The policy maker has a "target rate", x_t, and seeks to avoid deviations from this target. At the same time, the policy maker must also take into account that changing the buffer entails adjustment costs, which are proportional to a cost parameter λ. Assuming that both the cost of deviations from the target and the cost of changing the buffer are quadratic, the problem facing the policy maker is to minimize

$$(x_t - b_t)^2 + \lambda (b_t - b_{t-1})^2 . \tag{1}$$

Solving for the optimal buffer gives

$$b_t = \frac{\lambda}{1+\lambda} b_{t-1} + \frac{1}{1+\lambda} x_t. \tag{2}$$

The optimal buffer rate is a weighted average of the current rate and the target rate, with transaction costs determining the extent of the movement towards the target. An alternative formulation is to look at changes in the buffer,

$$b_t = \frac{1}{1+\lambda} (x_t - b_t) , \tag{3}$$

which highlights the partial nature of the adjustment towards the target.

To make the model more interesting I add further elements. First, I assume that policy is implemented with a lag and is aimed at a moving and possibly stochastic target. The buffer set in the current period will only take effect in the next period. The present value of deviating from the buffer is therefore $\beta E_t \left[(x_{t+1} - b_t)^2 \right]$, where β is a discounting parameter.

Second, the model is made dynamic by considering a multi-period setting. The target moves according to

$$\Delta x_{t+1} = -\phi x_t + \psi \Delta x_t + \epsilon_{t+1}, \tag{4}$$

where $\Delta x_t = x_t - x_{t-1}$. $\phi < 1$ can be interpreted as a mean-reversion parameter and $\psi > 0$ as a momentum parameter. When the target moves far away from its mean value, assumed to be zero, it starts being pulled back (mean reversion). Moreover, changes in the target in one direction tend to be followed by further changes in the same direction (momentum). Finally, ϵ_t is a zero-mean shock term.

Equation (4) is a second-order difference equation and, given appropriate parameters values, the mean reversion and momentum effects combine to describe a target moving in a non-explosive cyclical fashion.

The policy maker seeks to minimize all future costs of being far from target while at the same time taking adjustment costs into account

$$-\sum_{t=0}^{\infty}\left[\beta^{t+1}\left(x_{t+1}-b_t\right)^2+\lambda\beta^t\left(b_t-b_{t-1}\right)^2\right]. \tag{5}$$

The problem can also be formulated in value function notation as

$$V(b_{t-1},x_t,x_{t-1})=\underset{b_t}{max}\ \left\{-\lambda\left(b_t-b_{t-1}\right)^2+\beta E\left[-\left(x_{t+1}-b_t\right)^2+V(b_t,x_{t+1},x_t)\right]\right\}. \tag{6}$$

This notation makes clear the trade-offs facing the policy maker. The policy maker inherits a buffer from the last period and must take into account the current and past targets to predict the future target. Given these state variables, the policy maker must trade off the cost of changing the buffer, the cost of deviating from the cycle in the next period, and the effect of the current decision on actions in future periods.

In this setting, the optimal buffer can be expressed in closed form. There are two steps involved in proving this (see the appendix for details). First, it is shown that the problem of solving (6) subject to (4) belongs to the class of linear-quadratic dynamic programming problems. This implies that the value function is a quadratic function of the state variables. We can then split the state variables into two groups: those concerning the target (let $\iota_t^\top = \begin{pmatrix} x_t & x_{t-1} \end{pmatrix}$) and those concerning the buffer rate. The value function can then be

written as

$$V(\iota_t, b_{t-1}) = \frac{1}{2}\iota_t^\top P_{ii}\iota_t + b_{t-1}^\top P_{bi}\iota_t - \frac{1}{2}b_{t-1}^\top P_{bb}b_{t-1} + d_t \tag{7}$$

The P-coefficients can be described by a Ricatti equation, which one must normally solve numerically. However, in this particular case there is a closed-form solution. In the second part of the proof, it is shown that the expression for the value function can be written in a way which is mathematically equivalent to a problem studied by (Gârleanu and Pedersen, 2013) for which they identify a closed-form solution.

Two characterisations of the optimal buffer

As shown in the appendix, the optimal buffer can be written as

$$b_t = \frac{1}{1+\bar{\lambda}+P_{bb}}\left(\bar{\lambda}b_{t-1} + E_t[x_{t+1}] + P_{bi}\begin{pmatrix} E_t[x_{t+1}] \\ x_t \end{pmatrix}\right) \tag{8}$$

This characterisation shows that the current buffer level (b_t) can be expressed as a linear combination of the buffer level in place (b_{t-1}), the current state of the cycle (x_t), and the expected future value of the cycle $E\left[x_{t+1}\right]$. Moreover, P_{bb} is increasing in adjustment costs, λ, and $\frac{1}{1+\beta\lambda+P_{bb}}$ is therefore decreasing in the adjustment cost. This term can therefore be thought of as controlling the magnitude of how much the buffer is used. Figure 1 illustrates these mechanisms. It shows that, unless adjustment costs are very high, policy (i.e. the buffer) is ahead of the current state of the cycle. If adjustment costs are very high, the buffer does not fully track the cycle, but is more muted.

Another way of looking at the problem is to consider changes in the buffer rather than the level. Equation (7) should equal equation (6) evaluated at the optimal buffer level. Differentiating both equations w.r.t. b_{t-1} and applying the envelope theorem shows that changes in the buffer can be written as

$$b_t - b_{t-1} = \frac{1}{\lambda}\left(P_{bi}\iota_t - P_{bb}b_{t-1}\right). \tag{9}$$

Figure 1: A stylized depiction of how the optimal buffer varies throughout the cycle for different values of the adjustment cost parameter, λ. Unless adjustment costs are very large, the buffer will be ahead of the cycle.

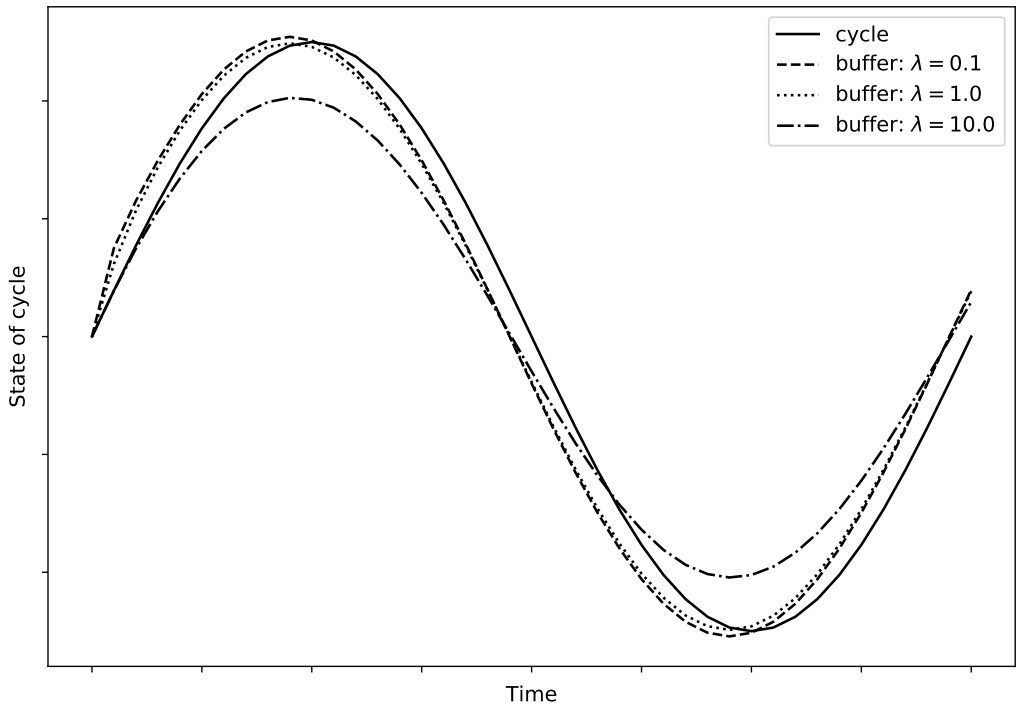

Figure 2: The figure shows the buffer dynamics in the case where the buffer starts out at a low value relative to the cycle (target). The dynamics are shown for different values of the adjustment cost parameter, λ. With low adjustment costs, there is a rapid catch-up.

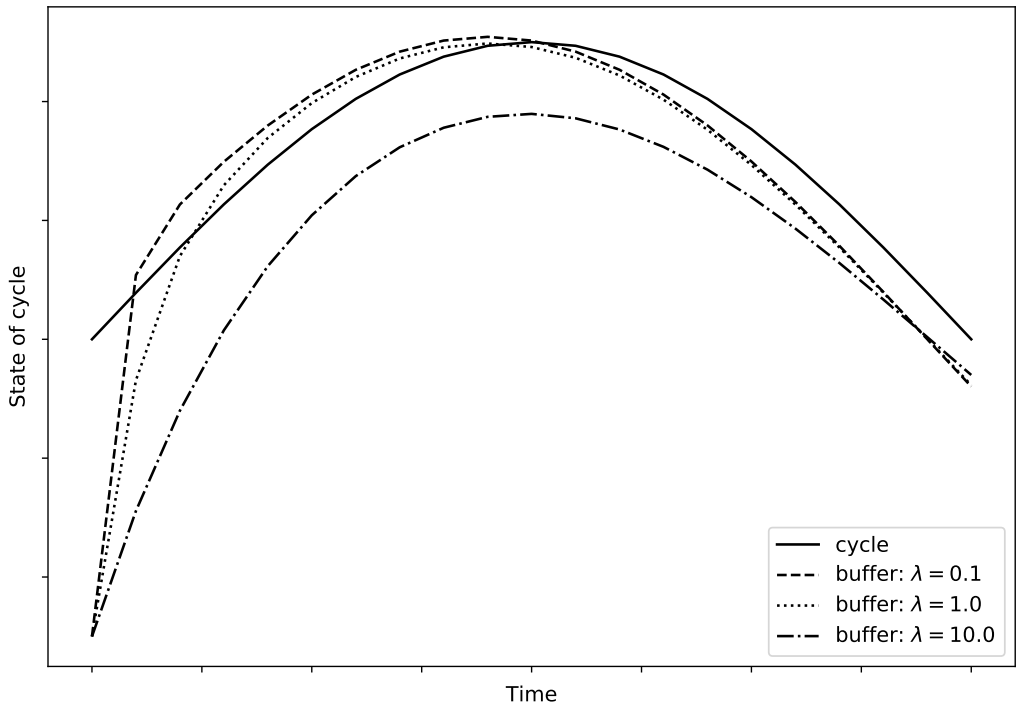

As a result, the optimal buffer can also be written as

$$b_t = \left(1 - \frac{P_{bb}}{\lambda}\right) b_{t-1} + \frac{P_{bi}}{\lambda} \iota_t \tag{10}$$

This simpler formulation describes the optimal buffer as a combination of the current buffer and the target state variables.

One can also look at the dynamics when the buffer starts out from a position which is out of place. Perhaps the most interesting case is that when the buffer is low relative to the current state of the cycle. In that case, the buffer rapidly catches up when adjustment costs are low, whereas the movement of the buffer is much more gradual when adjustment costs are larger. A stylized depiction of this is shown in figure 2.

3. General model

The above model can be viewed as a simpler, special case of a more general model. When deciding on the appropriate levels for policy instruments, macroprudential policy makers typically look at a variety of indicators and may need to consider multiple instruments and their interaction.

To fit this into a linear-quadratic setting, one can think of the policy maker identifying targets based on weighted combinations of indicators. As an example, when considering the countercyclical capital buffer policy makers assess a variety of indicators such as credit standards and developments, property prices, risks as judged by markets (e.g. credit spreads and volatility), model-based estimates of the state of the financial cycle, etc.[3] These indicators are frequently weighted and then illustrated in e.g. "heat maps" which show the signals coming from indicators grouped in various categories.

In this case we can write the targets (x_t) as linear combinations of the indicators (i_t)

$$x_t = W i_t \tag{11}$$

where x_t is a $k \times 1$ vector, W is a $k \times n$ matrix, and i_t is a $n \times 1$ vector.

In this more general setting, the changes in the targets come from the movement of the indicators

$$\Delta i_{t+1} = i_{t+1} - i_t = -\Phi i_t + \Psi \Delta i_t + \epsilon_{t+1}, \quad E_t\left[\epsilon_{t+1}\epsilon_{t+1}^\top\right] = \Sigma. \tag{12}$$

Finally, the cost of deviating from target and of adjusting the policy instruments ("buffers") are represented by $k \times k$ matrices Ω and Λ, respectively. The objective is therefore to maximize

$$-\frac{1}{2}E\left[\sum_{t=0}^{\infty}\left\{\beta^{t+1}\left(x_{t+1}-b_t\right)^\top \Omega\left(x_{t+1}-b_t\right) + \beta^t\left(b_t-b_{t-1}\right)^\top \Lambda\left(b_t-b_{t-1}\right)^\top\right\}\right] \tag{13}$$

[3] Appendix A in Danish Systemic Risk Council (2017) shows a concrete example of the set of indicators a macroprudential authority might look at.

subject to equations (11) and (12).

The general problem, likewise, has a closed-form solution (see the appendix), and one can characterise this solution in ways that are analogous to the way it was done for the simpler case. For example, the equivalent of equation (9) in the general case is

$$b_t - b_{t-1} = \Lambda^{-1} \left(P_{bi} \iota_t - P_{bb} b_{t-1} \right) \tag{14}$$

with P_{bb} and P_{bi} defined as in equations (A.20) and (A.21).

While the model is intended as a means for thinking qualitatively about macroprudential policy, it does, in principle, also offer a "recipe" for setting the levels of policy instruments. A macroprudential policy expert could use the model by deciding, using judgement, how to weight selected indicators to decide the level of the relevant policy instruments. A VAR-model could then be used to study the dynamics of these indicators, and, finally, the model could be applied based on estimates of how costly it is to depart from targets relative to the cost of adjusting behavior.

4. Adding realistic frictions: Simple numerical experiments

The world inhabited by actual macroprudential policy makers is more complicated than the setting studied in this paper. In this section, I consider a number of simple extensions to the model in order to identify how realistic frictions might affect the direction of macroprudential policy relative to the case where such frictions are not present. I study the frictions numerically and consider the case where the evolution of x_t is literally described by a cycle, which takes on values in the interval $[-2.5, 2.5]$.

One friction is that the policy implementation lag is typically not just "one period", but can span multiple quarters. How does changing the implementation lag affect policy? The numerical analysis shows that increasing the implementation lag simply moves policy forward, cf. figure 3. The longer the implementation lag, the further "ahead of the cycle" the

Figure 3: The figure shows the cycle and the buffer dynamics for implementation lags of 2, 4, and 6 periods respectively. The buffer value refers to the buffer value chosen in that period - which will then take effect either 2, 4, or 6 periods later. In all cases the adjustment cost parameter, $\lambda = 1$.

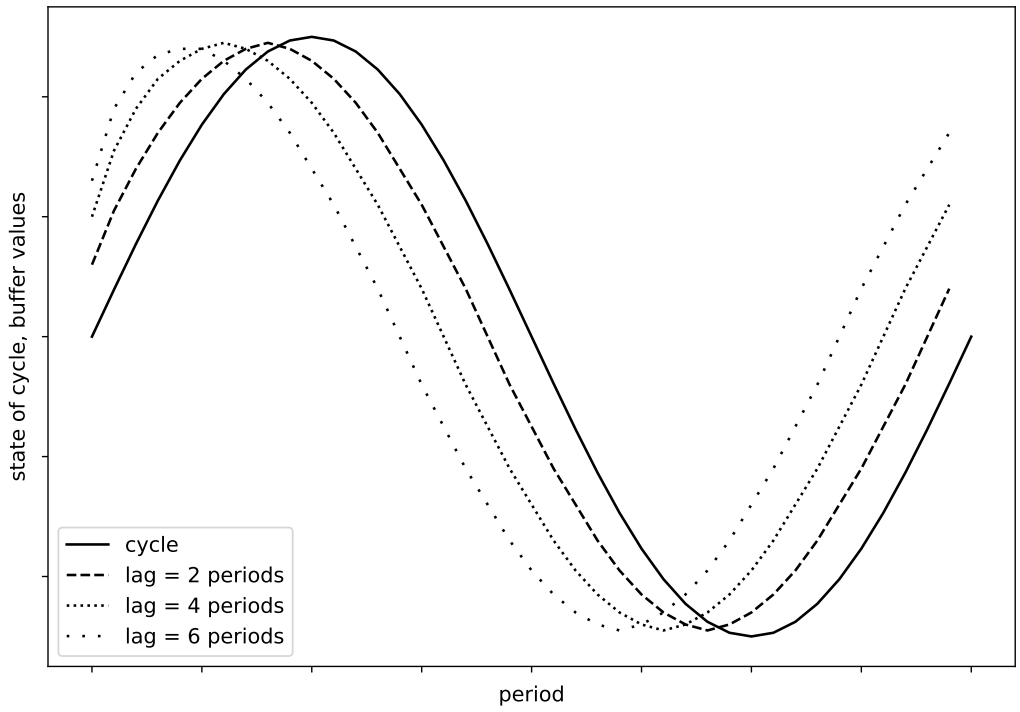

policy maker needs to be. Another friction is that macroprudential authorities have limited implementation means. An instrument such as the countercyclical buffer can only be chosen in increments of 0.25 and is ordinarily restricted to be between 0.0 and 2.5 per cent (under certain circumstances, though, a higher level is possible). Figure 4 shows the evolution of the buffer with a coarser grid of possible buffer values and restricting the buffer to positive values, relatively to the case of a relatively fine and unrestricted grid. Continuing from the above case where buffer values are restricted, I consider another extension, namely that the costs of changing buffers are likely to be asymmetric: It is costly to raise capital, but not to reduce it, and the cost of raising capital is presumably state-dependent, with raising capital being cheaper in upturns than downturns.

Figure 4: The figure shows the cycle and the buffer dynamics for different grids of possible buffer values. In all cases the adjustment cost parameter, $\lambda = 1$, while the policy lag is set to four periods.

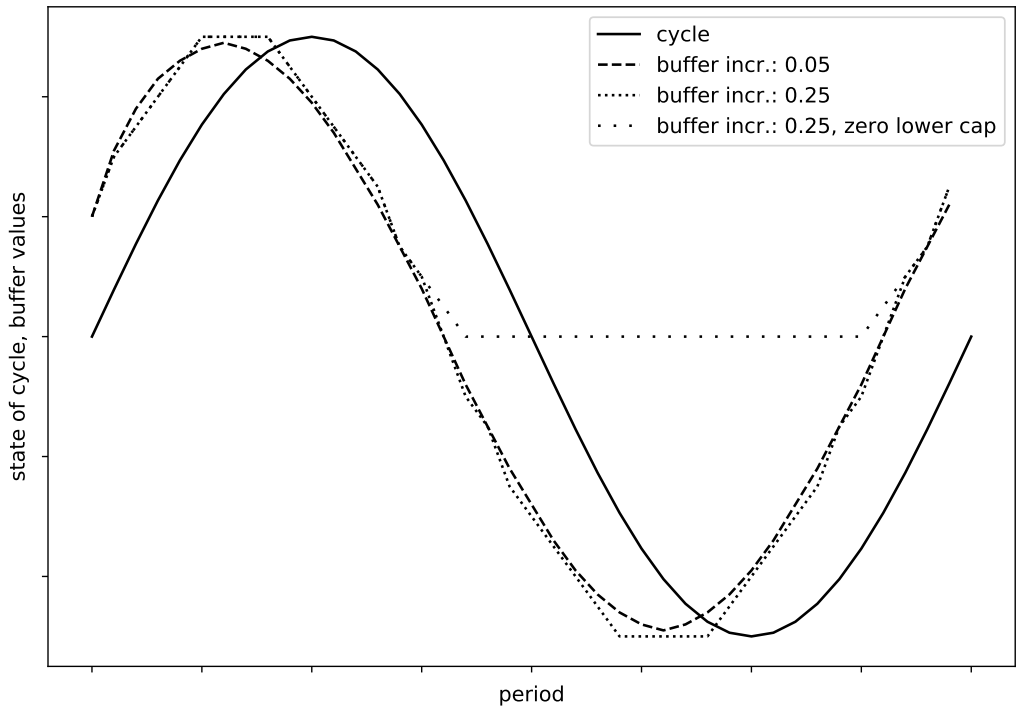

Figure 5: The figure shows the cycle and the buffer dynamics when buffer adjustment costs are (1) symmetric, (2) asymmetric (i.e. zero cost of reducing buffers), and (3) asymmetric and cyclical (lower cost in upturns). In all cases the adjustment cost parameter, $\lambda = 5$, while the policy lag is set to four periods.

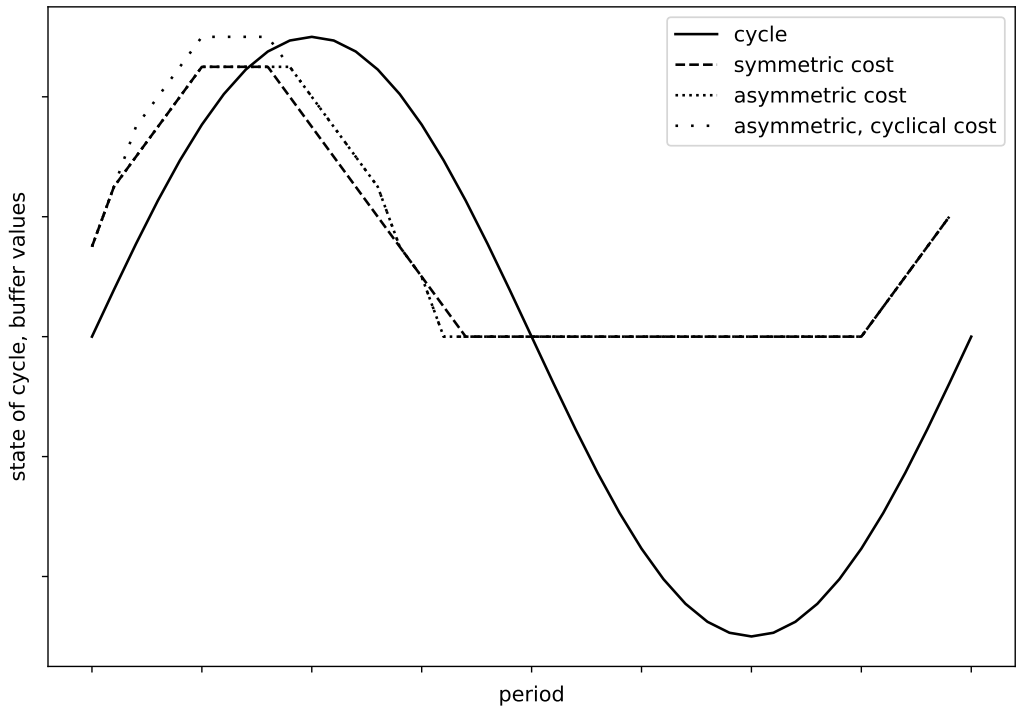

When costs are asymmetric, i.e. buffers can be reduced at no cost, the policy maker waits slightly longer to reduce buffers, but then reduces them faster, see figure 5. This better approximates the movement of the cycle than the buffer policy in the symmetric case. This is because, in the symmetric case, the policy maker tries to smoothe the reductions so they are of equal size in each period (which lowers adjustment costs). If costs are not only asymmetric, but also cyclical, there is a faster build-up of buffers and to a higher level.

The faster build-up of equity capital in good times appears consistent with the communication of policy makers who frequently advocate raising equity while it is cheap to do so.

5. Conclusion

When setting macroprudential policy, policy makers look at a variety of economic and financial indicators to gauge where the cycle is. This paper makes the point that policy makers should not only look at the present state of the cycle, but also where it is heading, and therefore need to be "ahead of the cycle" when setting e.g. the level of instruments such as the countercyclical capital buffer. The model thereby provides support and a rationale for the practical approaches to setting macroprudential policy followed by some central banks.

In terms of future research, one could consider other extensions to the model than those briefly mentioned in the paper. As an example, the paper treats banks as passive agents who immediately raise capital when capital buffers are increased. A more sophisticated analysis could look at the response of banks to the macroprudential regulation. Banks typically hold capital somewhat in excess of their capital requirements in order not to breach them. It is not obvious that banks, especially those mainly exposed to systemic risks, would have an incentive to further increase their capital in response to increased capital buffers if these buffers are withdrawn when systemic risks materialize. Another simplifying assumption of the paper is that it treats the cycle as being unaffected by the macroprudential policy, where at least some macroprudential instruments are potentially intended to dampen the cycle itself.[4] That is not the case for the countercyclical capital buffer, but a number of so-called "borrower-based" measures have this aim.

Finally, a more sophisticated analysis might take a closer look at the objectives of the macroprudential authority. The authority is likely to care mainly about robustness of policy and the most adverse states, such as financial crises, and setting policy instruments so as to mitigate the consequences in such states.

[4] A less sanguine possibility, as pointed out by Hórvath and Wagner (2017), is that macroprudential policy instruments can increase systemic risk taking because they insulate banks against aggregate shocks.

References

Bank of England, 2016. The financial policy committee's approach to setting the countercyclical capital buffer. Policy Statement April.

Bianchi, J., Mendoza, E.G., 2018. Optimal time-consistent macroprudential policy. Journal of Political Economy 126(2), 558–638.

Claessens, S., 2018. An overview of macroprudential policy tools. IMF Working Paper WP/14/214.

Danish Systemic Risk Council, 2017. The countercyclical capital buffer. Memo 17 November.

Drehmann, M., Borio, C., Gambacorta, L., Jimenéz, Trucharte, C., 2010. Countercyclical capital buffers: Exploring options. BIS Working Paper No. 317.

Drehmann, M., Borio, C., Tsatsaronis, K., 2011. Anchoring countercyclical capital buffers: The role of credit aggregates. International Journal of Central Banking December issue, 189–240.

ESRB, 2018. A review of macroprudential policy in the EU in 2017. European Systemic Risk Board April.

Farhi, E., Werning, I., 2016. A theory of macroprudential policies in the presence of nominal rigidities. Econometrica 84(5), 1645–1704.

Galati, G., Moessner, R., 2018. What do we know about the effects of macroprudential policy? Economica 85(140), 735–770.

Gârleanu, N., Pedersen, L.H., 2013. Dynamic trading with predictable returns and transaction costs. Journal of Finance 68(6), 2309–2340.

Hórvath, B., Wagner, W., 2017. The disturbing interaction between countercyclical capital requirements and systemic risk. Review of Finance 21(4), 1485–1511.

Jiménez, G., Ongena, S., Peydró, J.L., Saurina, J., 2017. Macroprudential policy, countercyclical bank capital buffers, and credit supply: Evidence from the Spanish dynamic provisioning experiments. Journal of Political Economy 125(6), 2126–2177.

Lim, C.H., Costa, A., Columba, F., Kongsamut, P., Otani, A., Saiyid, M., Wezel, T., Wu, X., 2011. Macroprudential policy: What instruments and how to use them? Lessons from country experiences. IMF Working Paper No. 11/238.

Ljungqvist, L., Sargent, Thomas, J., 2012. Recursive macroeconomic theory. The MIT Press.

.

Appendix: Proofs

General model

The objective is to maximize

$$-\frac{1}{2}E\left[\sum_{t=0}^{\infty}\left\{\beta^{t+1}\left(x_{t+1}-b_t\right)^{\top}\Omega\left(x_{t+1}-b_t\right)+\beta^{t}\left(b_t-b_{t-1}\right)^{\top}\Lambda\left(b_t-b_{t-1}\right)^{\top}\right\}\right] \tag{A.1}$$

That is, there are costs to changing buffers and it is costly if buffers stray from targets (both x_t and b_t are $k \times 1$ vectors).

And these costs are represented by $k \times k$ matrices Ω and Λ, both symmetric and positive definite.

The targets are linear combinations of various indicators, i.e.

$$x_t = W i_t \tag{A.2}$$

with W a $k \times n$ matrix and i_t a $n \times 1$ vector.

We assume that the indicators evolve according to

$$\Delta i_{t+1} = i_{t+1} - i_t = -\Phi i_t + \Psi \Delta i_t + \epsilon_{t+1}, \quad E_t\left[\epsilon_{t+1}\epsilon_{t+1}^{\top}\right] = \Sigma. \tag{A.3}$$

We want to prove two things:

1. That the model is linear-quadratic, thereby establishing existence and uniqueness of a solution.

2. That its solution can be written in a form that is equivalent to that in Gârleanu and Pedersen (2013), allowing us to write the solution in closed form.

Proof that the model is linear-quadratic

The generic linear-quadratic problem can be written as one of maximizing

$$-\sum_{t=0}^{\infty}\beta^t\left\{z_t^\top R z_t + u_t^\top Q u_t + 2u_t^\top H z_t\right\} \tag{A.4}$$

subject to

$$z_{t+1} = Az_t + Bu_t + C\epsilon_{t+1}. \tag{A.5}$$

where z_t are state variables and u_t are control variables.

In the model, we can treat i_t, i_{t-1}, and b_{t-1} as state variables and $b_t - b_{t-1}$ as a control variable. The law of motion can then be written as

$$z_{t+1} \equiv \begin{pmatrix} i_{t+1} \\ i_t \\ b_t \end{pmatrix} = \underbrace{\begin{pmatrix} I-\Phi+\Psi & -\Psi & 0 \\ I & 0 & 0 \\ 0 & 0 & I_b \end{pmatrix}}_{\text{"}A\text{"}} \underbrace{\begin{pmatrix} i_t \\ i_{t-1} \\ b_{t-1} \end{pmatrix}}_{z_t} + \underbrace{\begin{pmatrix} 0 \\ 0 \\ I_b \end{pmatrix}}_{\text{"}B\text{"}} \underbrace{(b_t - b_{t-1})}_{u_t} + \underbrace{\begin{pmatrix} I \\ 0 \\ 0 \end{pmatrix}}_{\text{"}C\text{"}} \epsilon_{t+1} \tag{A.6}$$

This is equivalent to the formulation of the law of motion in the generic model.

We next write the deviations from target, $x_{t+1} - b_t$, as

$$x_{t+1} - b_t = \underbrace{\begin{pmatrix} W & 0 & -1 \end{pmatrix}}_{\bar{W}} \begin{pmatrix} i_{t+1} \\ i_t \\ b_t \end{pmatrix} = \bar{W}\left(Az_t + Bu_t + C\epsilon_{t+1}\right) = \bar{A}z_t + \bar{B}u_t + \bar{C}\epsilon_{t+1} \tag{A.7}$$

Plugging this into the objective function yields

$$-\frac{1}{2}\sum_{t=0}^{\infty}\beta^t\left\{z_t^\top\left(\beta\bar{A}^\top\Omega\bar{A}\right)z_t^\top + u_t^\top\left(\Lambda + \beta\bar{B}^\top\Omega\bar{B}\right)u_t^\top + 2z_t^\top\left(\beta\bar{A}^\top\Omega\bar{B}\right)u_t^\top + c\right\} \tag{A.8}$$

which shows the equivalence with the generic problem. (Here c refers to terms that are irrelevant to the maximization problem.)

Solving the model in closed-form

Having established that the problem has a linear-quadratic structure, we know that the value function is of the form $V(z) = z^\top P z + P_0$. Here P_0 is included in the stochastic version of the problem. Because we know that the policy function is the same in the deterministic version of the problem, a feature known as the certainty equivalence principle, see Ljungqvist and Sargent (2012), we will proceed by analysing the problem without shocks.

For our purposes, it turns out to be convenient to separate the indicator (i_t and i_{t-1}) and buffer (b_{t-1}) state variables and write the value function as

$$V(\iota_t, b_{t-1}) = \frac{1}{2}\iota_t^\top P_{ii}\iota_t + b_{t-1}^\top P_{bi}\iota_t - \frac{1}{2}b_{t-1}^\top P_{bb}b_{t-1} \tag{A.9}$$

where $\iota_t = \begin{pmatrix} i_t & i_{t-1} \end{pmatrix}^\top$.

In value function notation, the problem can be written as

$$\begin{aligned} \beta^{-1}V(\iota_t, b_{t-1}) &= \max_{b_t} \left\{ -\frac{1}{2}(b_t - b_{t-1})^\top \bar{\Lambda}(b_t - b_{t-1}) \right. \\ &- \left. \frac{1}{2}(W_0\iota_{t+1} - b_t)^\top \Omega (W_0\iota_{t+1} - b_t) + V(\iota_{t+1}, b_t) \right\} \end{aligned}$$

where $\bar{\Lambda} = \beta^{-1}\Lambda$, $W_0 = \begin{pmatrix} W & 0 \end{pmatrix}$, and

$$\iota_{t+1} = \begin{pmatrix} I - \Phi + \Psi & -\Psi \\ I & 0 \end{pmatrix} \iota_t = \hat{A}\iota_t. \tag{A.10}$$

The right-hand side of the value function equation (i.e. the expression to be maximized) can then be written as

$$-\frac{1}{2}b_t^\top J b_t + b_t^\top j_t + d_t \tag{A.11}$$

with

$$J = \bar{\Lambda} + \Omega + P_{bb} \tag{A.12}$$

$$j_t = \bar{\Lambda} b_{t-1} + \left(\Omega W_0 + P_{bi}\right) \hat{A} \iota_t \tag{A.13}$$

$$d_t = -\frac{1}{2} b_{t-1}^{\top} \bar{\Lambda} b_{t-1} - \frac{1}{2} \left(W_0 \hat{A} \iota_t\right)^{\top} \Omega W_0 \hat{A} \iota_t + \frac{1}{2} \left(\hat{A} \iota_t\right)^{\top} P_{ii} \hat{A} \iota_t. \tag{A.14}$$

The policy rule, i.e. optimal choice of buffers is

$$b_t = J^{-1} j_t. \tag{A.15}$$

The policy rule is therefore a function of P_{bb} and P_{bi}.

The value function itself takes on value

$$\frac{1}{2} j_t^{\top} J^{-1} j_t + d_t. \tag{A.16}$$

We therefore have that

$$\beta^{-1} \left(\frac{1}{2} \iota_t^{\top} P_{ii} \iota_t + b_{t-1}^{\top} P_{bi} \iota_t - \frac{1}{2} b_{t-1}^{\top} P_{bb} b_{t-1}\right) = \frac{1}{2} j_t^{\top} J^{-1} j_t + d_t. \tag{A.17}$$

It is now possible to find the coefficients (P_{ii}, P_{bi}, and P_{bi}) by matching the left-hand and right-hand sides. This implies the following restrictions:

$$-\beta^{-1} P_{bb} = \bar{\Lambda} \left(\bar{\Lambda} + \Omega + P_{bb}\right)^{-1} \bar{\Lambda} - \bar{\Lambda}, \tag{A.18}$$

$$\beta^{-1} P_{bi} = \bar{\Lambda} \left(\bar{\Lambda} + \Omega + P_{bb}\right)^{-1} \left(\Omega W_0 \hat{A} + P_{bi} \hat{A}\right). \tag{A.19}$$

This formulation is analogous to that in Gârleanu and Pedersen (2013) (see equations (A7) and (A8) in their appendix), with the following mapping of notation: $A_{xx} = P_{bb}$, $A_{xf} = P_{bi}$, $\bar{\Lambda} = \bar{\Lambda}$, $\gamma\Sigma = \Omega$, $B = \Omega W_0 \hat{A}$, $I - \Phi = \hat{A}$, $\bar{\rho} = \beta$ and $\rho = 1 - \beta$.

We can therefore plug these expressions into equations A15 and A18, respectively, to get

the closed-form solutions.

$$\begin{aligned} P_{bb} &= \left(\beta\bar{\Lambda}^{\frac{1}{2}}\Omega\bar{\Lambda}^{\frac{1}{2}} + \frac{1}{4}\left(\left(1-\beta\right)^2\bar{\Lambda}^2 + 2\left(1-\beta\right)\bar{\Lambda}^{\frac{1}{2}}\Omega\bar{\Lambda}^{\frac{1}{2}} + \bar{\Lambda}^{\frac{1}{2}}\Omega\bar{\Lambda}^{-1}\Omega\bar{\Lambda}^{\frac{1}{2}}\right)\right)^{\frac{1}{2}} \\ &- \frac{1}{2}\left(\left(1-\beta\right)\bar{\Lambda} + \Omega\right) \end{aligned} \tag{A.20}$$

$$vec\left(P_{bi}\right) = \beta\left(I - \beta\hat{A}^{\top}\otimes\left(I - P_{bb}\Lambda^{-1}\right)\right)^{-1} vec\left(\left(I - P_{bb}\Lambda^{-1}\right)\Omega W_0\hat{A}\right). \tag{A.21}$$

Simpler model

To get at the closed-form solutions for setting the buffer in the simpler model, we first write down the mapping from the general model to the more complicated model. The buffer, b_t, is the same, except that it is a scalar rather than a vector. The adjustment cost is also a scalar with $\Lambda = \lambda$ and $\bar{\Lambda} = \bar{\lambda}$.

We let the cost of deviating from the target be one, i.e. $\Omega = 1$, so the adjustment cost is viewed relative to this cost.

Instead of letting the target be a weighted function of the economic indicators, we simply let the target itself follow the second-order difference equation (note that P_{bi} is now a 1×2 vector).

$$\iota_t = \begin{pmatrix} x_t \\ x_{t-1} \end{pmatrix}, \tag{A.22}$$

$$\hat{A} = \begin{pmatrix} 1-\phi+\psi & -\psi \\ 1 & 0 \end{pmatrix}, \tag{A.23}$$

$$W_0 = \begin{pmatrix} 1 & 0 \end{pmatrix}. \tag{A.24}$$

Inserting these expression into the formulas from the general model and simplifying gives

$$J = 1 + \bar{\lambda} + P_{bb} \tag{A.25}$$

$$j_t = \bar{\lambda}b_{t-1} + (W_0 + P_{bi})\hat{A}\iota_t \tag{A.26}$$

$$P_{bb} = \left(\lambda + \frac{1}{4}\left(1 + \frac{1-\beta}{\beta}\lambda\right)^2\right)^{\frac{1}{2}} - \frac{1}{2}\left(1 + \frac{1-\beta}{\beta}\lambda\right) \tag{A.27}$$

$$P_{bi}^{\top} = \beta\left(\frac{\lambda}{\lambda - P_{bb}}I - \beta\hat{A}^{\top}\right)^{-1}\left(W_0\hat{A}\right)^{\top} = \beta\left(\frac{\lambda}{\lambda - P_{bb}}I - \beta\hat{A}^{\top}\right)^{-1}\begin{pmatrix}1-\phi+\psi\\ -\psi\end{pmatrix} \tag{A.28}$$

We can now insert these expression to get the rule for the buffer:

$$b_t = \frac{1}{1+\bar{\lambda}+P_{bb}}\left(\bar{\lambda}b_{t-1} + \left(\begin{pmatrix}1-\phi+\psi & -\psi\end{pmatrix} + P_{bi}\hat{A}\right)\iota_t\right) \tag{A.29}$$

Noting that

$$\begin{pmatrix}1-\phi+\psi & -\psi\end{pmatrix}\iota_t = E_t[x_{t+1}] \tag{A.30}$$

and

$$\hat{A}\iota_t = \begin{pmatrix}E_t[x_{t+1}]\\ x_t\end{pmatrix} \tag{A.31}$$

we see that

$$b_t = \frac{1}{1+\bar{\lambda}+P_{bb}}\left(\bar{\lambda}b_{t-1} + E_t[x_{t+1}] + P_{bi}\begin{pmatrix}E_t[x_{t+1}]\\ x_t\end{pmatrix}\right) \tag{A.32}$$

DANMARKS NATIONALBANK
HAVNEGADE 5
DK-1093 COPENHAGEN K
WWW.NATIONALBANKEN.DK

DANMARKS NATIONALBANK