


Reports of the Advisory Scientific Committee

No 11 / December 2021

On the stance of macroprudential
policy

by
Stephen G. Cecchetti and Javier Suarez



ESRB

European Systemic Risk Board

European System of Financial Supervision

Contents

Executive summary	2
1 Introduction	5
2 A general framework for macroeconomic and macroprudential policy	8
Box A Assessing the stance of monetary policy	9
3 Macroprudential policy objectives and growth-at-risk	14
Box B Growth and banking crisis	16
4 Welfare foundations and a policy rule: an example	22
Box C Optimal policy in the CARA/normal case	23
Box D A graphical representation of optimal macroprudential policy	25
Box E The accuracy of the growth-at-risk-based welfare criteria under deviations from normality	30
5 Challenges in the implementation of macroprudential policy	35
6 Concluding remarks	39
References	40
Imprint and acknowledgements	44



Executive summary

Macroprudential policymakers are the risk managers of the financial system. Their role is to ensure that the probability and severity of a crisis is at a level which is consistent with the preferences of the citizens they serve. To fulfil this role successfully, the authorities require a framework consisting of a measurable goal, a set of tools, and a model linking the two. In the familiar case of monetary policy, the analysis of general economic and financial conditions, seen through a lens combining theoretical and empirical models with an agreed-upon objective, produces prescriptions for setting interest rates and adjusting the size and composition of central banks' balance sheets. Typically, a comprehensive framework delivers both a positive and a normative assessment of a policy stance. It allows policymakers and outside observers to evaluate whether the current settings are either too accommodative or too restrictive, both with regard to historical or theoretical norms and to levels which are considered appropriate to meet mandated goals.

Our aim in this report is to apply some of the lessons learned from the well-developed monetary policy framework to macroprudential policy. We hope to provide an alternative to the current predominantly narrative approach. With that in mind, we present an example of a simple formal model that allows us to draw clear conclusions as to the design of optimal macroprudential policies and the assessment of existing policy settings against such an optimal benchmark. This is challenging, given the fact that macroprudential policy is complex, with multiple intermediate objectives and an array of tools. Nevertheless, we believe it is worth trying to find an empirically implementable summary measure of macroprudential policy stance that can complement, rather than replace, the heatmaps and disaggregated risk assessments that characterise the multi-dimensional monitoring frameworks currently underpinning the work of the macroprudential authorities.

As an example of an empirically feasible approach, we discuss a case in which the ultimate goal of macroprudential policy is to minimise the frequency and severity of the economic losses arising from episodes of severe financial distress. To quantify this objective, we first take economic growth as a summary measure of the impact of economic performance on welfare. Then, based on empirical evidence, we argue that financial distress primarily influences the lower tail of the distribution of growth outcomes. This leads us to focus on the increasingly popular concept of growth-at-risk as a proxy for financial stability. Integrating growth-at-risk (or the related concept of growth-given-stress) into an optimal policy design problem allows us to deliver an empirically implementable prescriptive measure of macroprudential policy stance.

In its baseline formulation, the analytical model we present relies on a number of strong assumptions, which are: society's preferences for possible output growth outcomes have a specific form (i.e. they can be described by a utility function exhibiting constant absolute risk aversion); output growth is normally distributed; and policies have a linear impact on both the mean and the lower tail of output growth. Although this entire construction is quite specific, it provides a useful benchmark result as it implies that policymakers should seek to maintain a constant target distance between mean growth and a reference lower quantile of growth outcomes (e.g. the 10th percentile of the growth distribution). Furthermore, the optimal target distance is a decreasing function of



society's aversion to financial instability and of the relative effectiveness of the available policy tools in reducing the severity of low-tail growth outcomes without damaging mean growth.

As is the case with monetary policy, a prescriptive measure of macroprudential policy stance emerges from a comparison between the model-implied optimal policy and current conditions. In this case, if the forecast for the distance between mean growth and the reference lower quantile of growth deviates from the target, then there is a need to adjust the policy setting. While the empirical implementation of the approach will typically rely on generalisations and extensions of the baseline formulation, we are confident that the logic of its main implications, including the potential to provide a normative measure of macroprudential policy stance, will prevail.

The implementation of a comprehensive macroprudential policy framework faces numerous challenges. What exactly do the macroprudential authorities need to know and how might they obtain the appropriate information? Two issues concerning measurement accuracy loom large. First, given the sparsity of data, estimating the properties of the tail of the growth distribution (including the determinants that lead growth into such a tail) is a challenge. The level of difficulty associated with measuring the severity of low-tail events with precision should, in our view, guide the choice of the variable used to measure the growth implications of financial instability. We conclude that it might be advisable for policymakers to focus on the 10th percentile of the distribution of output growth (even if this reflects some adverse growth outcomes that are not necessarily triggered by episodes of financial distress) rather than the less likely (and less empirically knowable) 5th percentile.

Second, implementing any framework requires policymakers to have an accurate measure of the trade-offs they face. While tightening macroprudential instruments should reduce the probability and severity of crises, these benefits may come at the cost of lower trend growth. Due to data scarcity and the need to identify causal effects, inferring the existence and the size of this trade-off (on an instrument-by-instrument basis) is as difficult as it is critical. However, as we accumulate more experience, precision should increase, reducing the difficulty of the task.

To put our contribution into perspective, we note two issues. First, our discussion of the necessary elements of a theoretical and empirical framework that could lead to a measure of policy stance is based on a set of very stylised examples. There is no guarantee that the exact conclusions we draw will carry over to more complex, better articulated models of the economic and financial system. It seems likely, however, that any comprehensive macroprudential policy framework will rely on both the characterisation of the distribution of future economic outcomes relevant for a policymaker's ultimate target and an estimate of how the available policy tools influence that distribution.

Second, it is important to keep in mind that when the authorities reduce the likelihood of severely adverse outcomes, individuals – firms, households and investors – change the way they act in ways that could ultimately undermine the resilience of the system. Ironically, policies aimed at mitigating financial stress might sow the seeds of a future crisis. Our reduced-form treatment of the impact of macroprudential policy on economic outcomes does not explicitly address these instances of moral hazard. However, assuming the relevant causal effects of policy actions are properly estimated, it could account for the way in which these behavioural responses reduce the (net) effectiveness of specific policies.



To conclude, the aim of this report is to begin a discussion outlining the challenges faced by researchers and practitioners seeking to construct a macroprudential policy framework. In our view, while the endeavour will take time and will require contributions from various fields, there is every reason to believe that these efforts will ultimately improve the assessment, design and communication of macroprudential policy.



1 Introduction

Effective policy decisions emerge from careful deliberation and thoughtful analysis within a coherent framework. A carefully constructed quantitative and qualitative assessment lends focus to discussions between decision-makers, guides adjustments of instruments, provides for transparency in communication, and enhances accountability. In the familiar case of monetary policy, the analysis of general economic and financial conditions, seen through a lens combining theoretical and empirical models with an agreed-upon objective, produces prescriptions for setting interest rates and adjusting the size and composition of central banks' balance sheets. Typically, a comprehensive framework delivers a normative assessment of policy stance, allowing both decision-makers and observers to determine whether the current settings are either too accommodative or too restrictive to meet policymakers' mandated goals.

It is true that conventional monetary policy, with its generally univariate inflation objective and single interest rate tool, is far less complex than macroprudential policy. Nevertheless, we believe it is useful to start with a practical framework containing the same fundamental ingredients – an objective, a set of tools, and a model linking the two – with the aim of developing a measure of macroprudential policy stance. While it may seem uncharitable to say so, macroprudential policy is currently at the stage (if not worse) monetary policy was at more than half a century ago. In 1960, even though central banking was nearly three hundred years old and there were decades of information on prices, national income and employment, the monetary policy framework was much less developed and less structured than it is today.¹ As economists gradually refined monetary theory, eventually merging original Keynesian, monetarist and real business cycle elements into dynamic stochastic general equilibrium models, central bankers were able to construct a quantitative framework they could use to assess their policy stances. In parallel, academic contributions and institutional experience highlighted the benefits of independent governance structures for monetary policy.² Even so, the journey was agonisingly slow, and it took until the mid-1990s for a consensus to emerge.

Surveying the current landscape, we see that a majority of national and supranational jurisdictions have some type of macroprudential authority. Many were born out of the financial crisis of 2007-9, so are not much more than a decade old. Partly because this is such a recent enterprise, there is an active debate over how to formulate objectives, how to use the available tools, and how to structure governance. While the challenge is significant, we hope that the existing breadth of knowledge of economics and finance, as well as cooperation between academics and the authorities, will soon produce a consensus framework for guiding macroprudential policy decisions.

In producing this report, our aim is to move the process along, applying some of the lessons learned from the development of the agreed-upon monetary policy framework to the case of macroprudential policy. Putting this slightly differently, we hope to provide an alternative to the

¹ The Riksbank, founded in 1668, is the oldest central bank in the world. Central banking, however, is really a 20th century phenomenon – in 1900 there were only 18 central banks, by 2000 there were 173. See King (1999).

² There is an extensive literature on the benefits of central bank independence. See Bernanke (2010) for a survey and Dincer and Eichengreen (2014) for empirical evidence.



current predominantly narrative approach. With that in mind, we offer some examples as well as a perspective on how to measure the macroprudential policy stance in a more compact and systematic manner. While the analytical framework we propose is implementable (with a precision that will increase in line with the accumulation of modelling expertise, econometric techniques, data and experience), we see it as adding to, rather than replacing, the multi-dimensional monitoring framework currently used by the macroprudential authorities.

This report is divided into six sections, including this introduction and some concluding remarks. Section 2 describes a generic macroeconomic policy framework, and includes a discussion of the intrinsically normative notion of a policy stance. In Section 3, we begin by applying this logic to the case of macroprudential policy – the fact that macroprudential policy is complex, with multiple intermediate objectives and an array of tools, makes this challenging. Nevertheless, we believe it is important to aspire to identify an empirically implementable summary measure of macroprudential policy stance that complements the heatmaps and disaggregated risk assessments that currently guide what we see as a fragmented policy framework. As an example of what is currently feasible, first we take economic growth as a measure of welfare, and then we think of financial distress as shaping the lower tail of the distribution of growth outcomes. This leads us to use the increasingly popular concept of growth-at-risk (GaR) as a proxy for financial stability.

In Section 4 we present an example of a simple formal model that allows us to draw sharp conclusions with regard to the design of optimal macroprudential policies and the assessment of existing policy settings against such an optimal benchmark. In its baseline formulation, the model considers the analytically tractable case in which society's preferences in respect of possible output growth outcomes have a specific form (namely that they can be described by a utility function exhibiting constant absolute risk aversion), output growth is normally distributed, and the impact of policies on both the mean and the lower tail of output growth is linear. Although the latter construction is quite specific, it provides a useful benchmark result in which the policymaker seeks to maintain a constant target distance between mean growth and a reference lower quantile of growth outcomes (e.g. the 10th percentile of the growth distribution). As with monetary policy, a normative measure of macroprudential policy stance naturally emerges from a comparison of the model-implied steady-state policy with current conditions, so a policy will be deemed accommodative (restrictive) if the distance between mean growth and the reference low quantile exceeds (falls short of) the steady-state optimum. Also, as in the case of monetary policy, a need to adjust the policy setting emerges should the forecast for the distance deviate from the optimal target (which, in our simple reference case, happens to be constant and thus always equal to the steady-state target).³

Finally, we consider implementation in Section 5. Taking the growth-at-risk-based model results as a guide, we discuss how a policymaker might set a target in practice. What exactly do the macroprudential authorities need to know and how might they obtain this information? Two issues concerning measurement accuracy loom large. First, authorities need to be able to estimate the

³ With a state-contingent optimal target, the optimal policy might be accommodative or restrictive relative to the steady-state optimal policy and current policy settings might be referred to as “excessively accommodative”, “insufficiently accommodative”, “excessively restrictive” or “insufficiently restrictive” by making a further comparison between the current distance and the state-contingent optimal distance.



properties of the tail of the growth distribution; something made challenging by the limited availability of data. Ideally, we would have data for a long history – a period during which the economic and financial structure remains unchanged. Unfortunately, this is rarely the case. How difficult it is to measure the severity of low tail events (and the impact of policy on such events) with precision should, in our view, guide the choice of variable. We conclude that it might be advisable for policymakers to choose the 10th percentile of the distribution of output growth rather than the less likely (and less empirically known) 5th percentile.

Beyond the need to measure the objective itself, policymakers require accurate estimates of the elasticity of the distribution of the objective with respect to their macroprudential instruments. In our concrete application, where we take output growth to be a measure of welfare, this is the distribution of future growth outcomes. Specifically, in order to set the optimal target and also to adjust the available instruments, the authorities need to know what impact their actions will have on mean growth and the lower percentile of growth that they have chosen as their benchmark. This is exactly analogous to the requirement that, in pursuit of their stabilisation objective, monetary policymakers need numerical estimates of the impact of interest rate changes on prices and output or employment. In the case of macroprudential policy the task is made difficult by the sparsity of data, the heterogeneity of macroprudential tools and their impact across jurisdictions, and the challenges posed by policy endogeneity. On this final point, it is important to keep in mind that what matters for policy design is not the historical correlation between tools and the relevant economic outcomes, but the causal impact of the tools on such outcomes. This may entail going beyond current approaches that estimate the reduced-form impact of policy variables on mean growth and the lower quantile of growth, e.g. by building structural models that explicitly capture how macroprudential policy innovations influence such outcomes.




2 A general framework for macroeconomic and macroprudential policy

To develop a measure of policy stance, we begin with a general macroeconomic framework in which the economic system is characterised by a set of impulses amplified by a propagation mechanism, leading to economic outcomes. The impulses are a set of real sector shocks to productivity or the terms of trade; nominal shocks to the interest rate, exchange rates, or asset prices; and financial shocks including changes in risk attitudes or new information about institutions' exposure and solvency. The propagation mechanism is the structure of the economy and the financial system. The amplification of the shocks depends on a variety of factors, including the structure of household, firm, and bank balance sheets as well as financial markets and infrastructures. There are generally two types of outcome or goal: traditional macroeconomic stability, including stable growth, high employment and stable inflation; and financial stability, understood to be characterised by a low frequency and modest severity of breakdowns in the provision of essential financial services such as payments or credit.

Figure 1 lays out this generic framework. We make no attempt to be exhaustive in our description of the sources of impulses or the conditions which influence the strength or weakness of the propagation mechanism. Instead, we list the components of the system that are the most relevant for examining monetary and prudential policy.

Figure 1

A generic macroeconomic framework



The stability of the system, both macroeconomic and financial, depends on:

1. the dynamic stochastic properties of the shocks that hit the system;
2. the degree to which the various mechanisms amplify and propagate shocks over time and across agents, activities and markets.

Within this context, consider the familiar textbook case of conventional monetary policy – the policymakers' problem has three critical elements. First, express the objective in the form of a loss function to be minimised – for example, the weighted sum of squared deviations of inflation from its target and current output from potential output. Second, specify a policy tool, such as the short-term nominal interest rate. Third, postulate a model connecting the two, embedding a propagation mechanism that links shocks and current and future interest rate movements to inflation and output deviations. Importantly, we have a clear sense of the steady-state optimal or long-run equilibrium level of the policy interest rate, as well as an idea of how it should respond to shocks that push inflation and output away from their target levels. In Box A we present a stylised analytical model of the monetary policy problem.⁴

Looking at the generic framework, we generally cast the central bankers' problem as one where they work to meet their stabilisation objective by *reacting* to shocks which, if they were allowed to propagate, would destabilise the system. In other words, monetary policy interventions short-circuit, mitigate or neutralise the impact of otherwise harmful impulses on the targeted outcome.

Box A

Assessing the stance of monetary policy

Consider a simple model where the monetary policymaker controls the path of the interest rate i_t to minimise a quadratic loss function in inflation and output gaps, subject to a set of linear constraints that represent the dynamic path of the economy. We can write this as:

$$(A.1) \quad \min_{i_t} L = E[(\pi - \pi^*)^2 + \lambda(y - y^*)^2],$$

subject to

$$(A.2) \quad \begin{bmatrix} \pi_t - \pi_t^* \\ y_t - y_t^* \end{bmatrix} = A(L) \begin{bmatrix} \varepsilon_t^d \\ \varepsilon_t^s \\ (i_t - i_t^*) \end{bmatrix},$$

where π is inflation, y is output, ε^d is a demand shock that moves output and inflation in the same direction, ε^s is a supply shock that moves output and inflation in opposite directions, i is the nominal interest rate, and $A(L)$ is a matrix of lag polynomials in the lag operator L . In addition, π^* is target inflation, y^* is potential output and i^* is the neutral (or steady-state optimal) interest rate, i.e.

⁴ Our discussion of monetary policy is illustrative – we do not intend it to be a literal description of current practice. Over the last dozen years or so, because of the 2007-9 financial crisis and the COVID-19 pandemic, monetary policy has become significantly more complex, with the result that there are now numerous (potentially competing) objectives. In addition to price stability, central bankers act to stabilise the availability of funds to diverse sectors of the economy, as well as various segments of financial markets.



the interest rate that equals π^* plus the real interest rate consistent with output being equal to potential output.^(a)

There are three points to note. First, we can derive the loss function from microeconomic foundations by taking the second-order approximation to a stochastic intertemporal representative-agent utility maximisation problem in which there are costs associated with price adjustment (see Chapter 2 of Woodford (2003)). Second, the weight on output fluctuations in the loss function (λ) determines the speed at which the optimal policy returns inflation to its target (see Svensson (1999)). So, a relatively low λ is consistent with a central bank having a hierarchical mandate in which inflation is the primary objective. Third, since we can always characterise shocks in a more complex linear model as linear combinations of aggregate demand and aggregate supply shocks, the reduced-form description of the economy in (A.2) can accommodate more structural characterisations of the economy's dynamics.

Solving the problem in (A.1) and (A.2) gives rise to an interest rate path that depends on the history of the shocks. We can write the loss-minimising level of the policy rate ($\tilde{i}_t$) as:

$$(A.3) \quad (\tilde{i}_t - i_t^*) = b_1(L)\varepsilon_t^d + b_2(L)\varepsilon_t^s,$$

where $b_1(L)$ and $b_2(L)$ are lag polynomials with coefficients that are functions of the elements in $A(L)$ and the weight on output deviations in the loss function.^(b) When setting the interest rate, policymakers will seek to neutralise demand shocks – this is possible because interest rates move output and inflation in the same direction. For supply shocks there is an output-inflation volatility trade-off and policy will adjust depending on the parameter λ in the loss function – the smaller λ is, and the more significant inflation fluctuations are, the bigger the reaction will be.

Incidentally, we can always invert the dynamic structural system (A.2), writing it as a reduced form. Inserting this form into the optimal policy rule (A.3) yields a generalised version of the Taylor rule whereby deviations of the policy rate from its equilibrium level are a function of current and past deviations of inflation from its target, output from potential, and the policy rate itself.

Finally, turning to the concept of policy stance, when $\tilde{i}_t = i^*$, so that the current policy rate equals its steady-state optimal level, the stance is neutral. When the policy rate exceeds this neutral level, we refer to the stance as “restrictive”; when the interest rate is below this neutral level, we refer to the stance as “accommodative”. If a policymaker were to set policy some distance away from the optimal path, the qualifiers “excessively” or “insufficiently” could be added to the adjectives “accommodative” and “restrictive” to indicate the direction in which policy settings should adjust to be closer to the optimal path.

(a) It is straightforward to include a vector of exogenous variables in the dynamic system by simply adding a term from $B(L)X_t$ to the system (A.2). Through this generalisation, the optimal policy rule (A.3) would include reactions to the history of the X 's as well as to the ε 's.

(b) We could also write (A.3) as a partial adjustment expression where $\tilde{i}_t = i^*$ depends on its own lags, as well as the history of the shocks.



This monetary policy framework yields a natural measure of policy stance: the level of the interest rate relative to its steady-state optimal level (i^*). If the policy rate exceeds this level, policy is restrictive; if the policy rate is below the steady-state optimal level, policy is accommodative.⁵

It is certainly true that monetary policy and, more generally, the presence of an active central bank, can influence the structure of the financial system, the level of risk taking and, eventually, financial stability. However, as Svensson (2017) argues, in realistic model calibrations the inflation and output losses that would arise from using monetary policy as a financial stability tool outweigh the benefits by a factor of several hundred. Similarly, in an open economy new Keynesian dynamic stochastic general equilibrium model with financial frictions, Ozkan and Unsal (2014) conclude that when macroprudential policy tools are available there is no welfare gain from monetary policy reacting to credit growth. In a related paper, Ajello et al. (2016) show that even if the central bank were concerned about financial stability, its optimal monetary policy would react only very modestly to financial stability risks.⁶ In addition, the IMF (2015, Box 3) provides empirical evidence which shows that even though interest rate increases may have a positive effect on financial stability in the long run, they have a negative effect in the short run.

Turning to macroprudential policy, following Tucker (2015) we can frame the general role of financial stability policy as addressing a problem of “the commons” which is analogous to grazing on public lands or fishing in public waters.⁷ The “tragedy of the commons” arises when individuals have an incentive to do things that degrade the environment for everyone else. From this perspective, we can interpret financial stability as a common resource that is non-excludable yet rivalrous. If the financial system is stable, no one can be prevented from basking in the glow of its stability.

Importantly, individuals can act in ways that reduce systemic resilience. Just as a farmer has the incentive to overgraze, letting their cows eat until the public green becomes bare leading to the starvation of others’ herds and eventually their own, an actor in the financial system may have incentives to take risks that, because of spillovers, can deplete systemic resilience putting others at risk. Excessive risk-taking incentives may be exacerbated by the response of a financial firm’s owners and managers to the presence of both a social safety net (in the form of deposit insurance, the lender of last resort, and implicit government guarantees) and limited liability. When the risk taken by one agent affects outcomes for others, there is a classic externality: the insolvency of one

⁵ An alternative, explicitly prescriptive, measure of monetary policy stance compares the level of the interest rate with that implied by the optimal rule at each point in time. That is, minimising the objective, subject to the economy’s dynamic path, yields an optimal instrument rule. Using such a reference point, the stance measure would tell us whether policy is optimal, above optimal or below optimal, not just whether it is accommodative or restrictive. Combining the two criteria would allow us to describe policies as optimally neutral, accommodative or restrictive, as well as whether they are insufficiently or excessively accommodative or restrictive.

⁶ There is an active and fertile debate examining the relationship between monetary policy and prudential policy. For example, Cecchetti and Kohler (2014) provide a simple example in which interest rate and capital requirements are substitutes. Others, including Collard et al. (2017) and Mendicino et al. (2020), examine cases in which policies interact, so coordination could be beneficial.

⁷ See Cecchetti and Tucker (2016) for more details.



firm can cascade, creating system-wide runs, fire sales and an economy-wide credit crunch as balance sheets shrink.⁸

Policymakers can use their prudential toolkit to counter these externalities, pushing individual investors and institutions to internalise the costs their actions impose on others. The ESRB (2019) describes this as a process in which calibrating the tools requires policymakers to set their objective in the form of a “net systemic risk” (or “risk-resilience gap”) standard, monitor the level of risk and resilience in the system, and then adjust their policy stance to maintain the desired level of net systemic risk in the face of material changes to both the distribution of possible shocks and the fragility of the system.

In principle, financial stability policy and monetary policy are similar. In both cases a policymaker needs a well-defined and measurable goal, a set of tools, and models linking the two. For example, a macroprudential policymaker might focus on preventing acute system-wide disruptions to the provision of financial services that are essential for the proper functioning of the economy. System-wide disruptions in credit intermediation, liquidity and payment services, insurance, asset management, market-making services and the like are a characteristic feature of financial crises.

We now translate this relatively vague mandate to maintain the provision of financial services into an objective notion of what it means to pursue financial stability: acute disruptions of financial services should be infrequent and, when they do occur, the implications for the real economy should not be severely adverse.⁹ Given this goal of a low frequency and modest severity of system-wide disruptions, the macroprudential policymaker has a set of tools that might include changing the level of capital requirements, imposing maximum loan-to-value ratios for residential mortgages, modifying sectoral risk weights in capital requirements, and defining alternative stress test scenarios used to assess and influence the levels of resilience of relevant financial players, to mention just a few. In order to reach their goals, macroprudential policymakers must also have some idea of the conceptual and quantitative link between their tools and their mandated objectives.

In terms of the generic framework presented in Figure 1, we think of macroprudential policy as primarily influencing the propagation mechanism, maintaining financial stability by ensuring that the system remains *resilient* to shocks (e.g. by influencing the buffers through which different agents in the system may be able to absorb shocks). That said, the distribution of shocks likely depends on the state of the economy and the conditions in the financial system, and in particular agents’ risk-taking decisions that can, in turn, be shaped by policy. This endogeneity means that by reducing risk taking throughout the system macroprudential policy may also have an influence on the nature and size of the shocks affecting the system. To illustrate the point, consider the well-known case of booms and busts in property markets that may be caused by bubbles or simply by the evolution of

⁸ See Hanson, Kashyap and Stein (2011) for a detailed discussion of the externalities that are the basis for broad-based capital and liquidity regulation.

⁹ This interpretation of financial stability is consistent with the statutory mandate of the ESRB in Regulation (EU) 2019/2176 of the European Parliament and of the Council which reads: “The ESRB should contribute to preventing or mitigating systemic risks to financial stability in the Union and thereby to achieving the objectives of the internal market.” The regulation goes on to define term systemic risk as “a risk of disruption in the financial system with the potential to have serious negative consequences for the real economy of the Union or of one or more of its Member States and for the functioning of the internal market.”



beliefs. Real estate is often leveraged, so when property prices collapse the impact can cascade through the system. Those households that are unable to meet their mortgage payment obligations may cut back on other consumption purchases, reducing aggregate demand. Some borrowers may even default, risking damage to lenders. In this case, there is a potential for a bigger shock in the form of a property price collapse accompanied by balance sheet fragility, which leads to greater amplification. Policymakers could reinforce resilience to such shocks by, for instance, using tools that force agents to operate with lower leverage.

When and how macroprudential policymakers should utilise the instruments at their disposal, and with what intensity, are the key decisions they face. In the unlikely event that employing macroprudential tools entailed no costs, policymakers would face no trade-off. If they could reduce systemic risk without harming growth or any other relevant measure of social welfare, then maximum resilience would be the target. Unfortunately, however, the most stable financial systems are almost always either small and underdeveloped or repressed. So, while such systems present little risk to stability, they might provide insufficient support to economic wellbeing as measured by economic growth or any other suitable proxy for society's welfare. The stability we seek is not the stability of the graveyard.



3 Macprudential policy objectives and growth-at-risk

In order to apply the generic framework in Figure 1 to the case of macroprudential policy, the first step is to specify the objective. This is more complex than it is in the case of monetary policy, where there is a broad consensus as to the desirability of some form of flexible inflation targeting in which central bankers seek to minimise an average of squared deviations of inflation from its target and output from potential over a certain time horizon. By contrast, macroprudential policy currently follows a more disaggregated process in which authorities separate the assessment of risks, the design of associated tools, and the implementation of offsetting interventions into a set of categories explicitly linked to intermediate objectives.¹⁰ Current practice identifies the underlying sources of systemic risk arising from the actions of specific institutions or groups, and then fashions dedicated tools to address these risks. For example, regulators and supervisors use capital requirements to mitigate banks' solvency risk and loan-service-to-income limits to contain residential real estate risk. This piecemeal approach has a significant appeal. At a theoretical level, it is consistent with the absence of a comprehensive, integrated framework that incorporates all aspects of the financial system and the real economy, combining intermediate objectives and their associated tools into a single policy design problem. On practical grounds, current practice is coherent with the dispersion, in many jurisdictions, of the governance of macroprudential tools across authorities – each with their own narrow mandate.

Our aim is to explore the possibility of complementing this fragmented methodology with another containing a single unified goal for macroprudential policymakers. The logic of our analysis derives from the straightforward proposition that if each intermediate objective could be represented by a single variable, we could produce a solitary, measurable goal that aggregates all these objectives. Such a final objective should combine the welfare benefits of meeting each intermediate objective together with the potential welfare costs of using the available policy tools to influence the intermediate objectives, making it possible to consistently identify optimal macroprudential policy mixes.

While the advantages of having a measurable encompassing goal for macroprudential policy are clear, it is not at all obvious how to formulate such an overarching objective. The reason for this is that macroprudential policy has both aggregate and distributional effects, potentially influencing both the size and the growth of relevant macroeconomic variables such as output and consumption, as well as their distribution across states of nature, across sectors and within the population. While we are aware of these limitations, nevertheless, for the purposes of the

¹⁰ The strategy is clearly stated in Recommendation of the ESRB of 4 April 2013 on Intermediate Objectives and Instruments of Macro-prudential Policy (ESRB/2013/1), which states that “intermediate objectives should act as operational specifications to the ultimate objective of macro-prudential policy, which is to contribute to the safeguard of the financial system as a whole, including by strengthening the resilience of the financial system and decreasing the build-up of systemic risks, thereby ensuring a sustainable contribution of the financial sector to economic growth.” Besides this, it establishes that in terms of goals, the list of intermediate objectives “should include: (a) to mitigate and prevent excessive credit growth and leverage; (b) to mitigate and prevent excessive maturity mismatch and market illiquidity; (c) to limit direct and indirect exposure concentrations; (d) to limit the systemic impact of misaligned incentives with a view to reducing moral hazard; (e) to strengthen the resilience of financial infrastructures.”

remainder of this report we follow the path of those policymakers who focus on GDP growth as a summary measure of economic wellbeing. If, however, policymakers were to choose an alternative objective to account for additional important determinants of society's welfare, such as the distribution of income, the extent of carbon emissions, or any other feature not properly captured by GDP growth, then all we would have to change in the analytical framework presented below would be the definition of the variable representing the final objective.

Before turning to specifics, we should emphasise another important difference between monetary policy and macroprudential policy. At a practical level it is possible to change interest rates frequently and quickly, with an almost immediate impact. By contrast, it is not realistic to adjust macroprudential instruments such as capital requirements, position concentration limits and loan-to-value maxima from one day to the next. This could both delay and prolong the impact of these policies.

Importantly, while the impact of the instruments may be slow, we can still distinguish their steady-state calibration from their potential time variation. The case of Basel III capital requirements illustrates what we mean here. Regulators set a baseline minimum for the ratio of a bank's capital to its risk-weighted assets, while the structural characteristics of the financial system and the authorities' tolerance of the cost of banking crises determine the calibration of both the risk weights and the minimum.¹¹ On top of this minimum, the authorities have the option to set, among other add-ons, a time-varying countercyclical capital buffer (CCyB). Policymakers can adjust the CCyB to maintain resilience and prevent excess cyclicality in credit supply in the face of changes to economic and financial conditions. While the baseline settings of the instruments are critically important, the focus of our discussion is mainly on the time-varying dimension of macroprudential policies. Specifically, our interest is in measuring the settings of macroprudential policy tools relative to their optimal path in the medium term.

Turning to the distribution of output growth, Box B confirms that growth exhibits pronounced negative skewness. Very briefly, looking at information from 46 countries over the period 1960 to 2018, we see there were 97 banking crises. Of these, 13 resulted in three-year average growth that was more than two standard deviations below trend. Laeven and Valencia (2018) identify 151 banking crisis episodes in 119 countries over a period of 47 years. Of these, 83 were associated with output losses of more than 10% of one-year's GDP. At the other extreme of the distribution, while it is possible to find large positive growth rates in very long time series, these generally occur in the immediate aftermath of the physical destruction caused by a war. For example, while Italy grew by 35% and 19% in 1944 and 1945 respectively, since that time Italian annual growth has never exceeded 9%.¹²

¹¹ The Basel Committee on Banking Supervision (2010) provides the analysis used in the initial calibration of Basel III. Quantitative models addressing such a calibration more recently include Begenau and Langvoigt (2018), Mendicino et al. (2018) and Elenev, Landvoigt and Van Nieuwerburgh (2018).

¹² Adrian, Boyarchenko and Giannone (2019a) find negative skewness in US growth data, while Adrian et al. (2018) confirm this finding in a broader cross-section of countries.



Box B

Growth and banking crisis

To get a sense of the impact of banking crisis on growth, we merge the output per capita of the Maddison Project Database (Bolt and Luiten van Zanden, 2020) with Baron, Verner and Xiong's (2020) recently published banking crisis chronology. The full dataset, covering 46 countries from 1870 to 2018 and including 207 crisis episodes, is the longest and most comprehensive currently available. Using these data we compute the various characteristics of the non-overlapping three-year average growth rates in output per capita over two samples, one starting in 1870 and the other in 1960. To account for systematic country differences we normalise the data by subtracting each country's mean growth and dividing by its standard deviation (computed over the appropriate sample). In these data crises occur in 11% of three-year non-overlapping periods, which, assuming each crisis lasts for three years, is equivalent to a crisis starting in any given year with an unconditional probability of 3.7%, or once every 27 years.

Table B.1 reports various moments for these data. In order to characterise crisis episodes while allowing for flexibility in the timing of implied output losses, we take the minimum average growth for any of the three-year periods that include the years Baron, Verner and Xiong (2020) identify as the beginning of each banking crisis. Normalised per capita output growth exhibits several striking features. First, when we include both crisis and non-crisis years, the data exhibit negative skewness over both the full and the most recent sample periods. Second, a crisis results in three-year average growth that is roughly one standard deviation below the mean. To get a sense of what this implies, note that in the six decades since 1960, average growth in the median country was about 2.5%, with a standard deviation of about 2.25 percentage points. This means that during a typical banking crisis growth is roughly 0.25% on average for three years. Furthermore, our 1960-2018 sample includes 874 non-overlapping three-year periods (which works out at around 19 each for 46 countries). With 97 crises, this means that on average there is one crisis for every nine of these three-year intervals. Put slightly differently, once every 25 to 30 years there is a banking crisis that results in the loss of 7.5 percentage points of GDP.

Table B.1

Moments of normalised per capita output growth

	1870-2018		1960-2018	
	Full Sample	Crisis Episodes	Full Sample	Crisis Episodes
Mean	0.00	-0.74	0.00	-0.93
Skewness	-0.32	-0.94	-0.17	0.20
Excess Kurtosis	2.71	2.08	-0.01	0.34
10 th percentile	-1.11	-1.88	-1.24	-2.16
Number of observations	1872	207	874	97
Median across countries of average growth	1.99%		2.45%	
Median across countries of standard deviation of three-year growth	3.08%		2.27%	

Sources: Maddison Project Database (2020); Baron, Verner and Xiong (2020); and authors' calculations.

Notes: Data are deviations from the country mean of non-overlapping three-year average growth rates in standard deviation units. Countries are Argentina, Australia, Austria, Belgium, Brazil, Canada, Chile, Colombia, Czech Republic, Denmark, Egypt, Finland, France, Germany, Greece, Hong, Kong, Hungary, Iceland, India, Indonesia, Ireland, Israel, Italy, Japan, Korea, Luxembourg, Malaysia, Mexico, the Netherlands, New Zealand, Norway, Peru, the Philippines, Portugal, the Russian




Federation, Singapore, South Africa, Spain, Sweden, Switzerland, Taiwan, Thailand, Turkey, the United Kingdom, the United States and Venezuela.

Bad as this is, some banking crises are far worse. Since 1960 there have been 13 crises that resulted in three-year average growth of more than two standard deviations below the country mean growth. For a typical country, that means average growth of -2% for three years and a total loss of 13.5 percentage points of GDP.^(a)

The following graphs plot the distribution of normalised average three-year growth for the long and the short samples. In both cases the black lines display the smoothed frequencies of the three-year average per capita growth rates during normal (non-crisis) periods, while the red lines show the distribution of three-year average per capita growth rates during crisis periods. There are two points worth mentioning. First, as we would expect, crises are characterised by lower growth – the red lines are markedly to the left of the black ones. Second, the crisis distributions exhibit negative skewness and have more than one mode. The various modes seen during crises may reflect the existence of different types of banking crisis (distinguished by their varying degree of severity, due perhaps to the convolution of these crises with sovereign and currency crises).^(b)

Figure B.1
Distribution of normalised average three-year growth

(percentages)



Sources: See Table B.1.

Simple tabulations like these do not address the issue of causality. Without conducting a more detailed analysis we cannot infer the extent to which financial stress in general, and banking crises in particular, are responsible for low growth. Nevertheless, there does seem to be a connection. In a sample of 16 countries over the period 1980 to 2017, Aikman et al. (2019) find that two-thirds of the episodes in which growth is more than two standard deviations below the mean are preceded by credit booms. Furthermore, it is rare to see severely adverse downturns in aggregate activity in the absence of a crisis, just as it is rare to see robust growth during a crisis.

What these findings say about the connection between the distribution of per capita growth and episodes of financial stress (specifically banking crises) is consistent with other contributions that emphasise the link between financial instability and output losses. For example, looking at a set of

40 financial crises from 1980 to 2007, Cecchetti, Kohler and Upper (2009) compute that the average output loss in a crisis is nearly 20% of one year's GDP. Adrian, Boyarchenko and Giannoni (2019b) and IMF (2017) conclude that lower mean growth is associated with lower growth-at-risk (i.e. smaller values for the low quantiles of GDP growth). In addition, Falconio and Manganelli (2020) find that financial shocks have an asymmetric impact, increasing the negative skewness of the growth distribution.^(c)

(a) We note that the Laeven and Valencia (2018) dataset paints a more negative picture. Using data on over 120 countries from 1970-2017, the authors concluded that there were 151 systemic banking crises. More than half of these resulted in a cumulative output loss of more than 10% of one year's GDP.

(b) See Cecchetti, Kohler and Upper (2009) for a discussion of the similarities and differences between crises.


(c) We also note that business cycles tend to exhibit a similar asymmetry. See, for example, Kim and Nelson (1999) and Dupraz, Nakamura and Steinsson (2019).

Next, for the purposes of illustration, consider the stylised distribution of output growth shown in Figure 2. Where Y is the level of output or GDP, define $y_t = \ln(Y_t) - \ln(Y_{t-1})$ as the one-period growth rate of output and $f(y)$ the probability density function y_t . Label $\bar{y} = E(y)$ as the (positive) mean growth rate (or potential growth rate) of output. For the purposes of discussion, consider dividing the growth distribution into two disjoint intervals. The interval to the left of the (negative) level y^R contains severely adverse growth outcomes which we interpret as the typical result of the financial system being under stress or experiencing a crisis. The portion of the distribution to the right of y^R contains more benign growth outcomes which we interpret as most typical of normal, non-crisis times. The threshold y^R has a value-at-risk interpretation. If q is the probability of growth falling in the stress interval, then $y^R(q)$ is the growth-at-risk at this probability.¹³ For future reference, we also define growth-given-stress, $y^S(q)$, as the expected growth rate conditional on being below the threshold $y^R(q)$.

¹³ See Cecchetti (2008), and Adrian, Boyarchenko and Giannone (2019b).



Figure 2
Stylised probability density of output growth



We note that for a suitable choice of probability q the growth-at-risk threshold $y^R(q)$ need not separate crisis and non-crisis regimes precisely. For example, there could be severe business cycle downturns that do not qualify as financial crises in the left tail, as well as moderate financial stress episodes in which growth remains close to the mean and therefore remains in the unshaded portion of the distribution (as is the case for the two overlapping distributions in Figure B.1). However, measures of financial conditions and stress risk indicators are often constructed for the express purpose of signalling the probability and/or severity of poor growth outcomes over the next few years.¹⁴

To continue, we can define the distribution and chosen quantile for growth over any horizon in two ways. The first method considers a single period growth h -periods ahead: $y_{t+h} = \ln(Y_{t+h}) - \ln(Y_{t+h-1})$, while a second option focuses on the average growth over the next h periods: $y_{t,h} = (1/h)[\ln(Y_t) - \ln(Y_{t-1})]$. In both cases we can construct a density function over the quantity of interest and the corresponding values for both growth-at-risk and growth-given-stress.

There is an important difference between growth-at-risk and growth-given-stress. To illustrate this point, we construct a very simple example for a fixed probability of stress. Figure 3 shows distributions with two different shapes but with a 10% growth-at-risk equal to -4% in both cases. In other words, $y^R(10\%) = -4\%$ in both cases. The distribution in red is bunched near the threshold and has a correspondingly thin tail, while the distribution in blue is more spread out with a relatively fat tail. As a result, the growth-given-stress – the expected growth conditional on being below the threshold – is much more negative in the second distribution (the one in blue). For the red distribution $y^S(10\%) = -4.4\%$ and for the blue distribution $y^S(10\%) = -7.5\%$.¹⁵ Depending on their preferences, the authorities could be led to choose one formulation of the objective over the other.


¹⁴ See, for example, Hatzius, Hooper, Mishkin, Schoenholtz and Watson (2010) and Lang, Izzo, Fahr and Ruzicka (2019).

¹⁵ The example in Figure 3 is based on the Pareto distribution. We choose this because the tail of any arbitrary probability density converges to a Pareto distribution, or power function. See Reitano (2017), Section 6.2.1, p. 193.

In addition to the location and the shape of the lower tail, both of which could be related to the degree of resilience of the financial system, macroprudential policy is almost always concerned with long-run average growth. Importantly, there may be a trade-off between systemic resilience and growth. While a financial system that experiences frequent and deep financial crises is unlikely to support sustainable growth, some of the risk-taking behaviour that could lead to crises might increase the availability of funds to projects that raise growth and welfare in the long run.¹⁶ Moreover, in the same way that recessions allow high-productivity activities to replace marginal ones, some financial crises may have cleansing effects that reduce the possibility of an economy falling into a low-growth trap, mitigating the above-mentioned trade-off between mean growth and low tail growth.

Figure 3
Different probability densities for 10% growth-at-risk

(percentages)



Notes: Figure shows two Pareto distributions $f(y) = k_\alpha y^{-(\alpha+1)}$ with the same 10% growth-at-risk of -4%. The distributions differ in terms of the size of the parameter α and the (normalisation) constant k_α . For the blue-shaded distribution $\alpha = 2$ and for the red distribution $\alpha = 8$.

An advantage of a framework that relies on either growth-at-risk or growth-given-stress as proxies for financial stability is its potential to capture nonlinearities. In other words, it allows for the possibility that policy tools may have a differential impact on different parts of the distribution of the objective – whether this is growth, as in our example, or something else. To see how this takes place, note that standard empirical analyses in other policy fields, including monetary policy,

¹⁶ See Ranciere, Tornell and Westermann (2008).

estimate the elasticity of the mean of the policy objective, e.g. inflation, with respect to the policy instrument, e.g. an interest rate. This approach implicitly assumes that either policy actions simply shift the location of the distribution without changing its shape or that the impact on the shape of the distribution may be safely ignored. By contrast, quantile regression – the statistical method used to measure growth-at-risk – expressly allows for changes in the entire shape of the distribution (although analysts normally focus on just a few relevant quantiles). What this means is that a framework focusing on growth-at-risk can reveal whether policy, or any other conditioning variable including a measure of financial stress, has a differential impact on different parts of the distribution of the objective. In other words, the approach allows for both translations and deformations in the distribution of growth outcomes. This includes, but is not limited to, cases in which the economic and financial system can shift between regimes that might be more stable or less stable.



4 Welfare foundations and a policy rule: an example

The next step in formulating a measure of policy stance is to construct a model linking policymakers' tools to their agreed-upon objective. The discussion in the previous section leads us to conclude that either growth-at-risk or growth-given-stress might be good candidates for measuring the impact of financial instability on growth outcomes. Additionally, the macroprudential policymaker needs to be alert to the possibility of a trade-off in which actions that reduce the probability and severity of financial stress, raising growth-at-risk, may have a negative effect on average growth. Analogous to the inflation target in a monetary policy framework, a setup could be envisaged whereby elected officials provide the macroprudential authorities with a mandate based on striking an appropriate balance between improving growth-at-risk (y^R) or growth-given-stress (y^S) and damage to mean growth. For example, parliamentarians might instruct policymakers to focus on a given threshold probability and target some optimal distance between mean growth and either growth-at-risk (y^R) or growth-given-stress (y^S). Note that a hypothetical distance equal to zero that implies full stability might also imply very low mean growth and will therefore only be socially desirable if society is extremely averse to instability.

Suarez (2021) derives precisely this result for the case in which society's preferences for growth can be represented by a utility function exhibiting constant absolute risk aversion – growth is normally distributed and the macroprudential instrument has a negative linear impact on average growth and a positive linear impact on growth-at-risk. In this case, an optimal macroprudential policy keeps the gap between average medium-term growth and growth-at-risk constant at a certain target level. That is, $(\bar{y} - y^R)$ is set to a target level that depends on a combination of society's attitudes toward risk and the sensitivity of average growth and growth-at-risk in respect of the macroprudential instrument. Furthermore, when growth is normally distributed the gap between average growth and growth-given-stress is proportional to the gap between average growth and growth-at-risk, so we can express the constant target distance in terms of either quantity. Optimal policy also keeps $(\bar{y} - y^S)$ equal to a constant target – Box C provides more details.¹⁷

¹⁷ Suarez (2021) presents a static model with a single policy tool, thus abstracting from dynamics that may change the policy design problem in a number of important ways. This is especially true in the presence of multiple tools that have different time-series profiles in their impact on the distribution of growth. Two complications are worth noting. First, the optimal distance from mean growth to the growth-at-risk (or growth-given-stress) will likely be time-varying and will depend on the history of shocks to the economy. Second, the optimal path of the various tools will likely depend on a combination of such path of shocks and what may be complex intertemporal interactions between the tools.

Box C

Optimal policy in the CARA/normal case

Suarez (2021) examines a stylised one-period model in which the representative agent's preferences for output growth outcomes can be described by a constant-absolute-risk-aversion (CARA) utility function and growth rates are approximately normally distributed. As is well known, if an agent has CARA preferences over normally distributed outcomes, then their objective function may be expressed as the mean outcome less the agent's CARA coefficient multiplied by the variance of the outcomes. Using the fact that the distance between the mean and any quantile of the normal distribution is proportional to the standard deviation of the distribution, Suarez shows that the welfare of the agent (their expected utility) can be written, ignoring the horizon h , as

$$(C.1) \quad W = \bar{y} - \frac{1}{2} \omega [\bar{y} - y^R(q)]^2$$

where ω is a constant that is increasing in the risk aversion of the representative agent and decreasing in the probability q of the quantile to which growth-at-risk refers.^(a) So, welfare equals mean growth minus a term in the squared deviation of the q^{th} quantile from the mean.

To derive the optimal rule, Suarez assumes a linear structure: the mean and the q^{th} quantile of growth depend on a measure of systemic risk, R , and a macroprudential policy tool, τ :^(b)

$$(C.2) \quad \bar{y} = \alpha + \beta R - \gamma \tau,$$

and

$$(C.3) \quad y^R(q) = -\alpha_q - \beta_q R + \gamma_q \tau$$

where the α 's, γ 's and β_q are all positive and β can be positive or negative as long as it is greater than $-\beta_q$. The most important property of this system is that policy reduces mean growth while it raises the (negative) q^{th} quantile.^(c)

Maximising the quadratic objective (C.1), subject to (C.2) and (C.3), yields a rule in which policy is a linear function of systemic risk:

$$(C.4) \quad \tau = \phi_0 + \phi_1 R.$$

Furthermore, following this optimal rule implies keeping the distance between the mean and the q^{th} quantile constant:

$$(C.5) \quad [\bar{y} - y^R(q)] = \frac{1}{\omega} \left[1 + \frac{\gamma_q}{\gamma} \right]^{-1}.$$

Note that this constant optimal distance depends on two factors: the more risk averse society is, the higher ω is, and the smaller the optimal distance is; the more responsive to policy the q^{th} quantile is relative to the responsiveness of the mean (i.e. the bigger γ_q is relative to γ), the smaller the optimal distance is.

There are two points worth noting. First, in the case of the normal distribution the optimal distance from the mean to grown-given-stress (y^S) is proportional to the optimal distance from the mean to



growth-at-risk (y^R). As a result, we can substitute y^S for y^R in the analysis above, and all the results stand – the only change is that ω differs by a constant factor.

Second, as Suarez (2021) shows, it is a straightforward matter to generalise this example to allow τ to be a vector, so the policymaker has more than one tool. In this case tools can be ordered by the ratio of their impact on the q^{th} quantile to their impact on mean growth – the ratio of γ_q to γ for each tool. The most efficient tools are at the top of such a list. Furthermore, optimal policy should aim to keep $(\bar{y} - y^R)$ constant at the optimal distance implied by the most efficient tool.

(a) Suarez (2021), Appendix A1, derives the exact expression. For a coefficient of relative risk aversion ρ , and cumulative distribution functions of the standard normal $\Phi(\cdot)$, then $\omega = \rho / [\Phi^{-1}(q)]^2$. For example, when $q = 10\%$, $\Phi^{-1}(q) = -1.281$. So, for $\rho = 4$, $\omega = 3.12$. In addition, it should be noted that the symmetry of the normal distribution implies that in this model the objective function in (C.1) is the same for q and for $(1-q)$. Since the true distribution of growth is almost certainly negatively skewed, we view this as an analytical curiosity and focus on the illustrative value of the normal case as a benchmark case in which postulating an objective function like in (C.1) for $q < 0.5$ (i.e. with a focus on the lower tail) can be explicitly connected to the primitive preferences for growth outcomes of a risk-averse representative agent.

(b) This formulation abstracts from the case in which non-macroprudential policies have an impact on mean growth and growth-at-risk. One way to integrate such policies into the model is to reformulate the current measure for systemic risk, R , as a vector that includes these additional policies. They would then appear in a more general form of (C.2) and (C.3), as well as the macroprudential policy reaction function (C.4). In a more general discussion of optimal policy coordination, the framework might be further extended to cases in which the objective function W includes terms reflecting the goals of such policies. See Cecchetti and Kohler (2014) for an example that combines conventional monetary policy with capital regulation.

(c) A formulation in which policy influences some intermediate objective, which then alters the distribution of growth, is exactly equivalent. Specifically, Section 5.5 of Suarez (2021) also considers a case in which multiple intermediate objectives, each affected by targeted policy variables, have a non-linear effect on growth-at-risk, while policy still has a cost in terms of mean growth. In this case the optimal distance between mean growth and growth-at-risk is not constant but its determinants (and implied intuitions) are the same as in the formulation described here.

At this stage it is worth taking a moment out to discuss a key assumption leading to the conclusion that optimal policy targets the distance between mean growth and growth-at-risk, $(\bar{y} - y^R)$, i.e. that policies reducing the probability and/or severity of low growth outcomes (raising y^R) lower average growth ($\bar{y}$). This is a technical requirement in order to arrive at a nontrivial solution to the policy problem analysed in Suarez (2021). In the absence of such a trade-off, if policymakers had a tool that could raise growth-at-risk without lowering mean growth, the optimal policy would be to set policy to minimise the distance between the two. While such tools may exist, we strongly suspect that this is a local, rather than a global, property. This means that there may be a range over which the policy tool could both reduce the distance $(\bar{y} - y^R)$ and raise mean growth, but as the tool's setting increases, a trade-off will appear.¹⁸ Thus, we may view the linear equations of the model in Box C as an approximation to potentially non-linear relationships in the range over which policy entails a trade-off. (See Box D for a graphical representation of the policymaker's problem, both in the benchmark case where the stipulated linear trade-off is global and in a more general case where the slope of the elasticity of average growth with respect to macroprudential policy changes sign as the policy intensifies.)

Turning to the stance metric, we start by assuming that the policymaker's focus is on conditions h periods ahead. In other words, they perform what the inflation targeting literature calls “forecast

¹⁸ Looking at the model in Box C, this is a case in which the parameter γ in equation (C.2) is negative until τ reaches some critical level, at which point γ turns positive.



targeting” at horizon h . Since the influence of any policy changes takes time to work through the system, it is natural to target forecasts of future levels rather than current levels.¹⁹ Given the horizon, macroprudential policymakers will target the distance either from the mean to the growth-at-risk, $(\bar{y} - y^R)$, or from the mean to the growth-given-stress, $(\bar{y} - y^S)$. For the first of these we label the optimal target distance $(\bar{y} - y^R)^*$, and the stance then depends on the difference between $(\bar{y} - y^R)$ and $(\bar{y} - y^R)^*$. When the current expected difference is positive, $(\bar{y} - y^R)$ exceeds $(\bar{y} - y^R)^*$, policy is overly accommodative, and the tools need to be tightened. Conversely, if the expected difference is negative, policy is overly restrictive, and the tools need to be loosened.

Box D

A graphical representation of optimal macroprudential policy

To aid the reader’s understanding of the simple model described in Box C, we present the macroprudential policymaker’s problem in graphical form. We do this by first deriving a policy frontier and noting that optimal policy, if interior, is at the point where this frontier is tangent to the indifference curve arising from the welfare function.

Starting with the policy frontier, using equations (C.2) and (C.3) we can express the distance as a linear function of the measure of systemic risk R and mean growth $\bar{y}$:

$$(D.1) \quad [\bar{y} - y^R(q)] = \gamma_q \left[\frac{\alpha_q}{\gamma_q} - \frac{\alpha}{\gamma} \right] + \gamma_q \left[\frac{\beta_q}{\gamma_q} - \frac{\beta}{\gamma} \right] R + \left[1 + \frac{\gamma_q}{\gamma} \right] \bar{y}$$

(all parameters are defined in Box C).


This is a straight line with a slope equal to $1 + (\gamma_q/\gamma)$. We explore two cases in which the sign of the slope differs. When the slope is positive, macroprudential policy moves mean growth $(\bar{y})$ and the distance between mean growth and growth-at-risk $(\bar{y} - y^R)$ in the same direction, so there is a trade-off: reducing the distance comes at the cost of lower mean growth. We illustrate this case in Figure D.1 – the purple line is the policy frontier. To complete the graphical description of the model, we should note that the quadratic social welfare function (C.1) implies indifference curves in the depicted space such that, as mean growth increases and the distance decreases, welfare improves. These are the solid and dashed blue concave lines in the figure. Optimal policy, at the tangency of the policy frontier with the solid indifference line, is where the distance is equal to what we label d^* .

¹⁹ Svensson (1997) discusses this issue and its implications for policy design in a monetary policy setting.



Figure D.1

Optimal policy when there is a global trade-off



Sources: Based on Suarez (2021), Figure 1.


A property of the linear formulation in Suarez (2021) is that changes in the systemic risk variable R produce parallel shifts in the position of the policy frontier. Together with the shape of the indifference curves, this implies that the tangency points that would identify the optimal policy for each value of R always involve the same optimal distance d^* .

In the unlikely alternative case in which the slope of the policy frontier is globally negative, there is no trade-off. Policymakers could reduce the distance without there being any cost in terms of mean growth. In this case, the optimal policy is at the point where the downward sloping frontier cuts the horizontal axis.

An empirically more plausible case assumes the local absence of a trade-off at low levels of activation of the policy, followed by the emergence of a trade-off above a certain point. So, while the policy frontier might slope down over a certain range, it will eventually turn upwards. We illustrate a case of such a non-monotonic frontier in Figure D.2.

Figure D.2

Optimal policy when the trade-off only appears at high levels of policy activation



Sources: Based on Suarez (2021), Figure 1.

In this case there is still a tangency point that identifies the optimal policy and an interior optimal distance d^* associated with such a policy. However, in contrast to the baseline linear model of Suarez (2021) illustrated in Figure D.1, there is no guarantee that in this case the distance d^* will be invariant to changes in the systemic risk variable R . So, as postulated below, one can still use the distance associated with some steady-state level of systemic risk R^* as a benchmark for assessing the macroprudential policy stance. However, having an actual distance that is higher or lower than the implied d^* does not necessarily mean (if R is different from R^*) that such a policy is too loose or too tight.

The Suarez (2021) model suggests that the optimal distance $(\bar{y} - y^R)^*$ depends on three factors: i) the benchmark probability of stress (at the chosen horizon), ii) society's risk aversion, and iii) the impact of policy on the lower tail growth relative to its impact on mean growth (the quantity labelled γ_q/γ in Box C). The optimal distance increases as the probability declines, the risk aversion increases, or the relative impact goes down.

Figure 4 uses the exact expression Suarez derives (equation C.5) to compute the optimal target distance as the various determinants change. In the top panel we fix the threshold probability of

stress (q) at 10% and vary the coefficient of relative risk aversion (ρ) (which is a determinant of ω in equation C.5) from 2 to 6. The horizontal axis shows the relative impact of policy, while the vertical axis is the optimal target distance. When $\rho = 4$, and the relative impact (γ_q/γ), equals 5 (a value roughly consistent with the results reported in Galán (forthcoming)), the optimal target distance $(\bar{y} - y^R)^*$ is 6.84 percentage points. This number rises as risk aversion declines. Where $\rho = 2$, and the relative impact remains at 5, the optimal target distance rises to 13.69 percentage points.²⁰ In the bottom panel of Figure 4 we set the relative risk aversion (ρ) to 4 and vary the threshold probability q from 5% to 15%. Unsurprisingly, lowering the probability increases the distance. Focusing again on the case in which the impact of policy on growth-at-risk is five times as great as it is on long-run average growth, the optimal target distance falls from 11.27 percentage points at $q = 5\%$ to 4.48 percentage points at $q = 15\%$. The message we take from these very rough calculations is that for plausible parameterisations the optimal target distance implied by conventional relative risk aversion coefficients may be quite large – 10 percentage points or more. This suggests that unless policymakers are very averse to financial instability or have a macroprudential instrument that is extremely effective in improving growth-at-risk relative to its undesirable impact on mean growth (i.e. unless γ_q/γ is relatively large), using the policy tools to counteract the small probability of very large declines in output during crises may not be optimal. The case of snowstorms comes to mind. If we compare Stockholm, where it snows between 75 and 100 days per year, with Madrid, where a serious snowstorm occurs only once every half century, investing in snow removal infrastructure is a necessity in Stockholm, but not in Madrid.


²⁰ Estimates of the coefficient of relative risk aversion in the economic and financial literature vary depending on the methodology and reference data. Most economists would consider plausible values to be between 1 and 3, but empirical studies with microeconomic data yield estimates over a much wider range, depending on whether the focus is on consumption choices or portfolio decisions. For example, using subjective data on personal wellbeing, Gandelman and Hernandez-Murillo (2015) conclude that risk aversion is around 1 – a value close to that found in the public finance literature that they cite. By contrast, when trying to match asset prices and, specifically, the observed equity premium, the macro-financial literature requires much higher values. See, for example, Cecchetti, Lam and Mark (1994). In a macroprudential context, it is unclear whether the value of ρ should be set based on the preferences of a representative consumer, as in a typical macroeconomic model, or whether it should be inflated to reflect society's aversion to the cross-sectional unemployment and income inequality implications of financial crises.




Figure 4

Optimal target distance from mean growth ($\bar{y}$) to growth-at-risk (y^R)

A. Threshold probability of stress $q = 10\%$



B. Relative risk aversion $\rho = 4$



Sources: Authors' calculations based on equation (C.5) in Box C.

Returning to the issue of policy stance, recall that in the case of monetary policy we define a stance as restrictive or accommodative based on the level of the policy rate relative to its steady-state equilibrium level. Following this same line of reasoning, we posit that macroprudential policy is optimal when it maintains a target distance between mean growth and growth-at-risk (or growth-given-stress) that is consistent with the framework established above. Deviations from the optimal target distance $(\bar{y} - y^R)^*$ imply that a stance is either too tight or too loose. This means that, as is the case for monetary policy, we can evaluate the macroprudential policy stance by looking at the expected future path of relevant endogenous variables – in this case the central moment and the lower tail of the growth distribution.

We close this section with a discussion of the robustness of the results in Suarez (2021). Given the evidence showing that for most countries' growth is negatively skewed and has excess kurtosis, we are led to examine the accuracy of the normal approximation. As Box E shows, in all but the most extreme cases the constant distance rule in equation (C.5) performs quite well. Exceptions arise for low quantiles ($q = 5\%$) and high levels of risk aversion ($\rho = 20$).

Box E

The accuracy of the growth-at-risk-based welfare criteria under deviations from normality

If growth is normally distributed and utility exhibits CARA, then a rule maintaining a constant distance between average growth and a lower quantile of growth is optimal. This is the result obtained by Suarez (2021) that we highlight in our discussion. However, as we also emphasise, the empirical distribution of GDP growth is not symmetrical and may have fat tails. So, how much does this matter?

When the underlying growth distribution is not normal, we can think of a welfare metric that relies on moments of the growth distribution, such as the mean and the variance (as in classical portfolio theory) or the mean and a lower quantile of the growth distribution (as in the growth-at-risk framework), as an approximation for societal preferences. To judge the accuracy of this approximation when the distribution of growth is negatively skewed and has excess kurtosis, we examine a set of simulations in which growth is drawn from a mixture of two normal distributions (N_1 and N_2) with different means (μ_1 and μ_2) and standard deviation (σ_1 and σ_2). The two regimes capture “ordinary times” and “crisis times” respectively. Our assumptions are motivated by the distributions shown in Figure B.1. We should note, however, that (i) the distribution of growth in ordinary times has some bad outcomes but (ii) the distribution of growth in the event of a crisis has some relatively good outcomes. Put slightly differently, a crisis is neither necessary nor sufficient for growth to be low (or negative).

Our simulations vary across two dimensions: the likelihood of being in a crisis regime and the severity of a crisis (i.e. the mean of growth in the second regime). Table E.1 reports the moments of several illustrative distributions from our experiments. It shows cases for which the probability of being in a crisis regime over a three-year period is 5% and 10% and the average growth rate conditional on being in a crisis regime is 0%, -2% and -4%. In choosing these numbers, we are guided by the distributions for the period 1960 to 2018 reported in Table B.1, which suggest that a benchmark case – the highlighted row in the table below – is one in which about 10% of non-



overlapping three-year periods are crises and the average three-year annual growth rate during crises is 0%.

Table E.1
Candidate growth distributions

Average crisis growth (%)	Probability of being in a crisis regime							
	5%				10%			
	Mean	St. dev.	Skew	Excess kurtosis	Mean	St. dev.	Skew	Excess kurtosis
0.0	2.86	2.85	-0.10	0.14	2.71	2.70	-0.15	0.18
-2.0	2.76	2.76	-0.31	0.53	2.51	2.96	-0.42	0.54
-4.0	2.66	2.96	-0.61	1.27	2.31	3.30	-0.75	1.08

Notes: Unconditional moments of empirical distributions computed as the mixture of two normal distributions. The first has a mean of 3.01% and standard deviation of 2.52%; the second has a mean equal to the average crisis growth in the first column and a standard deviation of 2.78%. The probability of drawing from the second distribution is equal to 5% or 10%. All reported numbers are based on 500,000 draws. The shaded values are those that correspond to the benchmark in the data reported in Table B.1.

We proceed to draw 500,000 times from each distribution. From these draws, we first compute expected utility for a CARA utility function, which we can use to calculate the certainty equivalent growth rate. Then, using the result from Suarez (2021), we calculate the level of welfare implied by several alternative moment-based approximations (which are directly interpretable as approximations to this certainty equivalent growth rate). These are a simple mean-variance welfare criterion and welfare criteria based on the mean and the (squared) distance between the mean and 5%, 10% and 15% growth at risk (equation C.1 in Box C).


Figure E.1 reports the results for experiments in which the probability of being in the crisis regime varies from 0% on the left to 20% on the right. For these experiments we set average crisis regime growth at -2% and the coefficient of relative risk aversion at 4. The dashed line shows the deviation of certainty equivalent growth from the unconditional mean growth – what we think of as the representative agent’s “willingness to pay” to completely stabilise growth. It should be noted that these numbers are relatively small. For the case where the probability of being in the crisis regime is 10%, and so mean growth is 2.51% with a standard deviation of 2.96%, someone with a risk aversion of 4 is willing to pay 18 basis points of GDP to completely stabilise growth.^(a)

Turning to the accuracy of the growth-at-risk-based welfare criteria, the bars in the graph show the approximation error (deviation from the true certainty equivalent level) implied by each alternative welfare criterion. Two things stand out. First, the mean-variance criterion performs the best – something that is true in all the cases we explored. Second, except for the 5% growth-at-risk (GaR) rule (the orange bars), the errors are less than 2.5 basis points in absolute values. Putting this another way, in all but the worst-case scenario the errors are less than 10% of the willingness to pay to fully stabilise GDP growth.^(b)



Figure E.1

The accuracy of the growth-at-risk approximation varying the probability of a crisis regime




Notes: Results are based on 500,000 draws from a mixture of two normal distributions: $N_1(\mu=3.01, \sigma=2.52)$ and $N_2(\mu=-2.0, \sigma=2.78)$. The probability of drawing from N_2 is given by the value on the horizontal axis and relative risk aversion is set to 4.

Turning to a second set of experiments, we fix the probability of being in the crisis regime at 10% and vary the severity by varying the average growth in the crisis regime from +0.5% to -4.0% (Figure E.2). Everything else is the same as before: we rely on 500,000 draws from the mixture of two normal distributions and set relative risk aversion to 4. Willingness to pay remains modest at between 13 and 26 basis points. When crises are, on average, very severe, the errors implied by some of the growth-at-risk-based welfare measures now exceed 5 basis points. Having said that, the 10% growth-at-risk rule still performs remarkably well – the absolute errors remain below 1 basis point.

Figure E.2

The accuracy of the growth-at-risk approximation varying the severity of crisis regime



Notes: Results are based on 500,000 draws from a mixture of two normal distributions: $N_1(\mu=3.01, \sigma=2.52)$ and $N_2(\mu=\text{value on horizontal axis}, \sigma=2.78)$. The probability of drawing from N_2 is set equal to 10% and relative risk aversion is set equal to 4.


Next we report the results of a set of experiments in which relative risk aversion equals 20. In this case, as crises become sufficiently severe, willingness to pay (measured as the difference between certainty equivalent growth and unconditional average growth) rises above 100 basis points. The results are shown in Figure E.3. (Please note that the vertical scale extends much higher than it does in the previous two figures.) Perhaps unsurprisingly, in such an extreme case approximation errors can exceed 20 basis points, although they remain small as a fraction of willingness to pay.

Finally, we note that the results here are robust to a variety of alternative departures from normality. These include cases where the second (crisis) regime is characterised by either a Student's t-distribution with varying degrees of freedom (which has fatter tails than normal) or an exponential distribution (implying an unconditional distribution with sizeable negative skewness). The results are also robust (again in the sense that they produce relatively small approximation errors) when preferences feature constant relative risk aversion (CRRA) rather than CARA.

Overall, we find the results of these experiments reassuring. The message is that the errors arising from building policy design using welfare metrics that rely on the (squared) distance between mean growth and a lower growth quantile (especially the 10th percentile) are not only consistent with expected utility maximisation under CARA preferences and normality but provide a reasonably good approximation to the latter when the growth distribution deviates substantially from normality, as well as under CRRA preferences.^(c)

Figure E.3

The accuracy of the growth-at-risk approximation with high relative risk aversion



Notes: Results are based on 500,000 draws from a mixture of two normal distributions: $N_1(\mu=3.01, \sigma=2.52)$ and $N_2(\mu=\text{value on horizontal axis}, \sigma=2.78)$. The probability of drawing from N_2 is set to 10% and relative risk aversion is set to 20.

- (a) This result is reminiscent of Lucas's (1987) observation that, in a representative agent model, the cost of business cycles is likely to be extremely small. See Barlevy (2004) for a discussion of Lucas's original results and a survey of the related work.
- (b) An alternative way to evaluate the importance of deviations from normality is to compare the errors to revisions in published quarterly GDP growth numbers. Looking at the US case from the initial to the most recent estimate (from 1990 to 2015), the median GDP growth revision is more than 25 basis points – over five times the largest error in Figure E.1.
- (c) If growth were truly normally distributed, then a growth-at-risk approach would have no advantage relative to a more conventional mean-variance approach. The growth-at-risk approach can only be strictly superior if the conditional distribution of growth is not normal and, as recent evidence suggests, financial factors and macroprudential instruments have a differential impact on the lower quantiles of the distribution. In other words, when changes in systemic risk and policy settings do not simply shift the mean and change the variance of the distribution, they change its shape as well.

5 Challenges in the implementation of macroprudential policy

Policy design is an inherently empirical exercise. While we need conceptual models to discipline our thinking and ensure logical consistency, most policy actions involve quantities. Monetary policymakers set policy rates at certain levels, decide on the size and composition of their balance sheet, and so on. Prudential authorities are no different. Microprudential regulators set rules that establish minimal or maximal values for key ratios associated with the operation of individual financial intermediaries. Similarly, the macroprudential policy toolkit contains many quantitative instruments. Determining the appropriate stance requires measurement, evaluation and the calculation of an optimal policy response.

To see how we can proceed with measuring stance, take the case of the European Central Bank's (ECB's) monetary policy framework as a guide. As of May 2021, the ECB states its objective as price stability, which is defined as inflation (as measured by the year-on-year increase in the Harmonised Index of Consumer Prices for the euro area) of below but close to 2% over the medium term.²¹ This involves three essential elements: an index for measuring inflation, a horizon over which to measure it, and a specific number for the target itself. Once these are established, the Governing Council then assesses the policy stance based on whether its tools are set at levels most likely to meet the objective.

Applying this logic to the specific macroprudential policy framework we proposed in the previous sections of this report, there are three categories of input feeding into the construction of the optimal target distance between mean growth and downside risk that provides the benchmark for measuring stance. These are:

- (i) the index, horizon, and degree of time averaging;
- (ii) the threshold lower quantile and the choice of growth-at-risk or growth-given-stress;
- (iii) the effectiveness of policy, i.e. the impact of policy on the lower tail of output growth relative to mean growth (γ_q/γ).

We now consider the three categories of necessary inputs from both a conceptual and an empirical perspective. That means we discuss what we *should* measure as well as what we *can* measure.

i) The index, the horizon, and the degree of time-averaging

Starting with the index, we should choose an indicator that is closely tied to the general welfare of the society in question. In practice this means focusing on (the growth of) GDP, consumption or employment. The work done so far focuses primarily on the first of these, but we should not rule out alternatives.

²¹ At the time of writing the ECB is conducting a review of its monetary policy strategy, so this may change.

Turning to the horizons, we can justify looking forward four, eight, twelve or even sixteen quarters ahead. The choice depends in part on the lag with which policy influences financial risks. For example, changes in the countercyclical capital buffer (CCyB) may have to be announced with a lead time of four quarters and may take an additional four quarters to have any impact. In such a case it makes no sense for the objective to be at a shorter horizon than that required to implement the policy and for it to have any impact. In practical terms, the choice of horizon depends on the precision with which we can measure the impact of other required inputs on the target.

Regarding the degree of time-averaging, policymakers should decide whether to frame their objective in terms of a one-year growth rate h years ahead or the average growth rate over the next h years. In our view, the latter would be more natural.²² The rationale for this choice is that average growth takes account of the fact that the costs and benefits of macroprudential policies are almost certainly spread differently over time. To illustrate this point, consider a policy of tightening the maximum loan-to-value ratio requirement for residential mortgages. This could reduce expected growth one and two years out while reducing downside risks three and four years out. In such a case it makes sense to choose an objective based on average growth over the next three or four years. Importantly, such a measure implies less focus on short-lived fluctuations and more on low-frequency, persistent risks.

ii) The threshold lower quantile and the choice between growth-at-risk and growth-given-stress

Next, consider the choice of quantile and the characterisation of the lower tail of the growth distribution. Starting with the former, should macroprudential policy focus on the 5th percentile of the distribution or, possibly, the 10th or the 15th? At a conceptual level it is reasonable to consider lower quantiles. The Laeven and Valencia (2018) database implies an unconditional probability of a crisis of roughly 4.5% per year, suggesting that we should focus on the 5th percentile of the growth distribution. However, this seems too low for two reasons. First, financial factors play a role in most downturns – even downturns that are not accompanied by financial crises. Second, we suspect that there are significant barriers to measuring low quantiles with precision. As the quantile declines from the tenth to the fifth to the first, observations around the true quantile are very likely to become increasingly sparse, so the accuracy with which the quantile (and its determinants) can be estimated inevitably declines. In all, this is an argument for preferring the 10th quantile to the 5th (and also to the 15th, which might less clearly reflect the implications of financial stress).

Turning to the measure of the lower tail of growth outcomes: which option is best, growth-at-risk or growth-given-stress? The discussion in Section 3 of this report, as well as the example in Figure 4, suggests that choosing the latter might make more sense. Since growth-given-stress can vary substantially for a fixed growth-at-risk, and our concern is with extremely negative growth outcomes, it would be logical to focus on the expected shortfall, i.e. the growth conditional on the

²² For the sake of simplicity and ease of presentation, the framework we describe here abstracts from dynamics within the specified policy horizon and uses aggregation over such a horizon as a substitute for being explicit about the higher frequency path of the relevant state variables. Detailed articulation of the framework could instead rely on quantile vector auto-regressive models that explicitly capture such dynamics. Such a further evolution of the framework could also take account of (properly discounted) intertemporal trade-offs over the policy horizon (e.g. balancing short-term costs against what may be the medium-term benefits of a policy tool).

system being under stress.²³ However, there is a strong empirical case for choosing growth-at-risk. Computing growth-given-stress requires us to estimate the area under the entire lower tail, and the absence of data to pin down the density at very low quantiles would make this extremely difficult to do with any degree of precision. We cannot measure the frequency or the severity of events we very rarely see. So, much as we might prefer growth-given-stress as a measure of welfare, it seems prudent for policymakers to pay more attention to growth-at-risk.

iii) The relative effectiveness of policy

The final input into the computation of the macroprudential target is the impact of policy on the lower tail of the growth distribution relative to its impact on mean growth, (γ_q/γ) . This requires policymakers to estimate the elasticity of average growth for the chosen low quantile in respect of the array of macroprudential tools over the preferred horizon. Several complex issues arise in this regard. First, the accuracy of these estimates will almost certainly depend on the horizon. This means we will be able to estimate the impact of policy on growth more precisely at some horizons than at others – a fact that will play a role in the choice of the horizon itself. Second, we have more experience of some tools than others. For example, changes in maximum loan-to-value ratios for residential mortgages have historically been more common than adjustments to the CCyB or changes in bank asset concentration limits. If a tool shows no variation this means that available data will be silent on its effectiveness. Third, as we discussed in and around Box D in the previous section, there is a possibility that this trade-off may not apply to all settings of each policy tool. Finally, there is the issue of the endogeneity of policy tools. An appropriate treatment of macroprudential instruments' endogeneity is essential if estimates of (γ_p/γ) are to capture the causal effect of policy on the relevant moments of the growth distribution rather than the mere historical correlation between tools and growth outcomes.²⁴

These inputs, combined with society's aversion to severely adverse events (the coefficient of relative risk aversion ρ in the analysis in the previous section), provide a measure of the optimal target distance that is the basis for a macroprudential target. Comparing this optimal target with the distance implied by current policy settings yields a measure of stance. When the current estimate of the distance exceeds the optimal target, policy is too accommodative; when the current estimate of the distance is smaller than the optimal target, policy is too restrictive.

Before we conclude, we note several additional challenges that macroprudential policymakers face during implementation. First, there is the sheer number of tools available. Alam et al. (2019) tabulate 17 separate categories of macroprudential tools. Ideally, we would determine which tools are substitutes and which are complements, so that we can employ such tools in the best possible

²³ To see why this is the case, consider computing utility, conditional on it being in the tail of the distribution. Assuming we can approximate the utility function as a finite-order polynomial, then expected utility (conditional on $y < y^R$) is a function of the moments of the distribution describing the lower tail of growth outcomes. For the special case of the Pareto distribution shown in Figure 4, these are all functions of a single parameter that determines the shape of the tail. In other words, there is a class of utility functions and a distribution of growth outcomes for which welfare could be expressed as a function of growth-given-stress but not as a function of growth-at-risk.

²⁴ Addressing this issue may require moving beyond standard reduced-form quantile regressions by adopting either an instrumental-variables approach or a structural approach that explicitly models policy as an endogenous variable in a multi-equational system.

combinations, equating their marginal effectiveness.²⁵ Second, as always, policymakers need to avoid reacting to “noise”. Given how underdeveloped data systems are for some parts of the financial system (especially non-bank intermediaries), this is a particular risk. A related call for caution emerges when we recognise the potential for misspecification and estimation error that could plague the empirical models underpinning the kind of policy calculations envisaged above.²⁶ Third, as should be clear from our discussion, the policy target is likely to differ across jurisdictions. Attitudes to risk (or society’s aversion to financial instability) will diverge, as will financial structure and the effectiveness of different policy instruments. So, in a multijurisdictional area such as the European Union, providing a cross-country assessment of policy stance will involve the challenge of treating or accommodating country heterogeneity along some of the dimensions identified above (risk attitudes, effectiveness of available policy tools, etc.).

²⁵ See Suarez (2021), Section 5.4 for a general discussion of this problem.

²⁶ Such problems plague many aspects of both public and private decision-making. See, for example, Svensson and Woodford (2003) for a general discussion, Orphanides (2001 and 2003) for an examination of the impact of “noisy” information on monetary policy, and Jorion (1985) for a study of the problem in the context of international portfolio diversification.



6 Concluding remarks

The role of macroprudential policymakers is to ensure that the probability and severity of a crisis is at a level that is consistent with the preferences of the citizens they serve. To fulfil this task successfully they require a measurable objective, a set of tools that can influence their target, and a model linking the two. The problem is analogous to that faced by monetary policymakers as they strive to achieve price stability. Using this as a guide, this report presents an example of a framework in which optimal macroprudential policy requires policymakers to target the distance between average growth and a low quantile of growth. This distance depends on society's aversion to crisis and the degree to which tools can influence the mean and the lower tail of the growth distribution. Our example yields a normative measure of stance, which tells us whether macroprudential policy is excessively accommodative or restrictive.

Before we conclude it is important that we provide a few warnings. First and foremost, the purpose of this report is to provide a perspective on the problems faced by macroprudential policymakers. We discuss the necessary elements of a theoretical and empirical framework that could form a basis for constructing a measure of policy stance. We present stylised examples based on a simple model. There is no guarantee that the conclusions we draw will lead to more complex, better articulated models of the economic and financial system. However, it seems likely that a fully articulated macroprudential policy framework will include a horizon for the target, a measure of the lower quantiles of a suitable aggregate indicator of economic wellbeing (possibly GDP growth), and an estimate of the causal effect of the relevant policy tools on that distribution. A combination of data sparsity and the difficulty faced by policymakers in identifying the causal impact of macroprudential tools on their target makes this a challenging task.

Second, our simplified treatment of macroprudential policy abstracts from a well-known danger that plagues all stabilisation policy. When the authorities reduce the likelihood of severely adverse outcomes, people change their attitudes toward risk taking in ways that could ultimately make the system less resilient. Ironically, policies aimed at mitigating financial stress could sow the seeds of future crises. Some elements of crisis management, in which authorities rescue financial markets and institutions, may further aggravate this problem. Our treatment of the impact of macroprudential policy on systemic risk (proxied by its impact on the low tail of the growth distribution in our example) does not account for this form of moral hazard. That said, if the moral hazard effects were dominant in practice, a suitably estimated measure of the causal impact of policy actions on the relevant low tail of the growth distribution would reflect this by showing an overall negative, rather than positive, effect of crisis mitigation policies on tail outcomes, and the framework envisaged in this report would advise against such policy actions.

To conclude, the goal of this report is to begin a discussion, outlining the challenges that researchers and practitioners face as they set out to construct a macroprudential policy framework. In our view, making progress on the road ahead will take time and will require contributions from various fields, but there is every reason to believe that these efforts will help to improve the assessment, design and communication of macroprudential policy.



References

- Adrian, Tobias, Boyarchenko, Nina and Giannone, Domenico (2019a), "**Vulnerable growth**", *American Economic Review*, 109(4), pp. 1263-1289.
- Adrian, Tobias, Boyarchenko, Nina and Giannone, Domenico (2019b), "**Multimodality in macro-financial dynamics**", *Staff Reports 903*, Federal Reserve Bank of New York.
- Adrian, Tobias, Grinberg, Federico, Liang, Nellie and Malik, Sheherya (2018), "**The term structure of growth-at-risk**", *CEPR Discussion Papers*, No 13349.
- Aikman, David, Bridges, Jonathan, Hacıoglu Hoke, Sinem, O'Neill, Cian and Raja, Akash (2019), "**Credit, capital and crises: a GDP-at-risk approach**", *Bank of England Staff Working Papers*, No 724.
- Ajello, Andrea, Laubach, Thomas, López-Salido, David and Nakata, Taisuke (2016), "**Financial stability and optimal interest-rate policy**", *Federal Reserve Board Finance and Economics Discussion Series*, 2016-067.
- Alam, Zohair, Alter, Adrian, Eiseman, Jesse, Gelos, Gaston, Kang, Heedon, Narita, Machiko, Nier, Erland and Wang, Naixi (2019), "**Digging deeper – Evidence on the effects of macroprudential policies from a new database**", *IMF Working Papers*, No 19/66, International Monetary Fund, Washington, DC.
- Barlevy, Gadi (2004), "**The cost of business cycles and the benefits of stabilization: A survey**", *NBER Working Papers*, No 10926.
- Baron, Matthew, Verner, Emil and Xiong, Wei (2020), "**Banking crises without panics**", *NBER Working Papers*, No 26908, March.
- Basel Committee on Banking Supervision (2010), "**An assessment of the long-term economic impact of stronger capital and liquidity requirements**", *Bank for International Settlements*, Basel, August.
- Begenau, Juliane and Langvoigt, Tim (2018), "**Financial regulation in a quantitative model of the modern banking system**", *Stanford University Business School Working Papers*, No 3558.
- Bernanke, Ben S. (2010), "**Central bank independence, transparency, and accountability**", *speech at the Institute for Monetary and Economic Studies International Conference*, Bank of Japan, Tokyo, May 25.
- Bolt, Jutta and Luiten van Zanden, Jan (2020), "**Maddison style estimates of the evolution of the world economy. A new 2020 update**", *Maddison Project Database, version 2020*.
- Cecchetti, Stephen G. (2008), "**Measuring the macroeconomic risks posed by asset price booms**", in Campbell, J.Y. (ed.), *Asset Prices and Monetary Policy, Chicago, III*, University of Chicago Press for NBER, pp. 9-43.



Cecchetti, Stephen G. and Kohler, Marion (2014), “**When capital adequacy and interest rate policy are substitutes (and when they are not)**”, *International Journal of Central Banking*, No 10(3), pp. 205-232.

Cecchetti, Stephen G., Kohler, Marion and Upper, Christian (2009), “Financial crises and economic activity”, *Proceedings - Economic Policy Symposium - Jackson Hole*, Federal Reserve Bank of Kansas City, pp. 89-135.

Cecchetti, Stephen G., Lam, Pok-sang and Mark, Nelson C. (1993), “**The equity premium and the risk-free rate: Matching the moments**”, *Journal of Monetary Economics*, No 31(1), pp. 21-45.

Cecchetti, Stephen G. and Tucker, Paul M. W. (2015), “**Is there macroprudential policy without international cooperation?**”, in Glick, R. and Spiegel, M. (eds.), *Policy Challenges in a Diverging Global Economy, Proceedings of the Asia Pacific Policy Conference*, Federal Reserve Bank of San Francisco, November.

Collard, Fabrice, Dellas, Harris, Diba, Behzad and Loisel, Olivier (2017), “**Optimal monetary and prudential policies**”, *American Economic Journal: Macroeconomics*, No 9(1), pp. 40-87.

Dincer, N. Nergiz and Eichengreen, Barry (2014), “**Central bank transparency and independence: Updates and new measures**”, *International Journal of Central Banking*, No 10(1), pp. 189-253.

Dupraz, Stéphane, Nakamura, Emi and Steinsson, Jón (2019), “**A plucking model of business cycles**”, *NBER Working Papers*, No 26351, October.

Elenev, Vadim, Landvoigt, Tim and Van Nieuwerburgh, Stijn (2018), “**A macroeconomic model with financially constrained producers and intermediaries**”, *NBER Working Papers*, No 24757.

European Systemic Risk Board (2019), “**Features of a macroprudential stance: initial considerations**”, *ESRB report*, April.

Falconio, Andrea and Manganelli, Simone (2020), “**Financial conditions, business cycle fluctuations and growth at risk**”, *Working Paper Series*, No 2470, European Central Bank, Frankfurt am Main.

Galán, Jorge E. (forthcoming), “**The benefits are at the tail: uncovering the impact of macroprudential policy on GDP growth**”, *Journal of Financial Stability*.

Gandelman, Nestor and Hernandez-Murillo, Ruben (2015), “**Risk aversion at the country level**”, *Review, Federal Reserve Bank of St. Louis*, No 97(1), pp. 53-66.

Hanson, Samuel G., Kashyap, Anil K. and Stein, Jeremy C. (2011), “**A macroprudential approach to financial regulation**”, *Journal of Economic Perspectives*, No 25(1), pp. 3-28.

Hatzius, Jan, Hooper, Peter, Mishkin, Frederic S., Schoenholtz, Kermit L. and Watson, Mark W. (2010), “**Financial conditions indexes: A fresh look after the financial crisis**”, *U.S. Monetary Policy Forum*.



International Monetary Fund (2015), “**Monetary policy and financial stability**”, *IMF Staff report*, August 28.

Jorion, Philippe (1985), “**International portfolio diversification with estimation risk**”, *Journal of Business*, No 58(3), pp. 259-278.

Kim, Chang-Jin and Nelson, Charles R. (1999), “**Friedman's plucking model of business fluctuations: Tests and estimates of permanent and transitory components**”, *Journal of Money, Credit and Banking*, No 31(3), pp. 317-334.

Lang, Jan Hannes, Izzo, Cosimo, Fahr, Stephan and Ruzicka, Josef (2019), “**Anticipating the bust: a new cyclical systemic risk indicator to assess the likelihood and severity of a financial crisis**”, *Occasional Paper Series*, No 219, European Central Bank, Frankfurt am Main, February.

Laeven, Luc and Valencia, Fabian (2018), “**Systemic banking crises revisited**”, *IMF Working Papers*, No 18/206, 14 September.

Lucas, Robert E. Jr. (1987), *Models of Business Cycles*, Basil Blackwell, New York.

Mendicino, Caterina, Nikolov, Kalin, Suarez, Javier and Supera, Dominik (2018), “**Optimal dynamic capital requirements**”, *Journal of Money, Credit and Banking*, No 50(6), pp. 1271-1297.

Mendicino, Caterina, Nikolov, Kalin, Suarez, Javier and Supera, Dominik (2020), “**Bank capital in the short and in the long run**”, *Journal of Monetary Economics*, No 115, pp. 64-79.

Ozkan, F. Gulcin and Unsal, Filiz D. (2014), “**On the use of monetary and macroprudential policies for small open economies**”, *IMF Working Papers*, No 14/112.

Orphanides, Athanasios (2001), “**Monetary policy rules based on real-time data**”, *American Economic Review*, No 91(4), pp. 964-985.

Orphanides, Athanasios (2003), “**Monetary policy evaluation with noisy information**”, *Journal of Monetary Economics*, No 50(3), pp. 605-631.

Ranci re, Romain, Tornell, Aaron and Westerman, Frank (2008), “**Systemic crises and growth**”, *Quarterly Journal of Economics*, No 123(1), pp. 359-406.

Reitano, Robert (2017), “**Foundations of Quantitative Finance**”, *Book 4: Distribution Functions and Expectations*, Brandeis International Business School.

Suarez, Javier (2021), “**Growth-at-risk and macroprudential policy design**”, ESRB Occasional Paper Series No. 19, September.

Svensson, Lars E. O. (1997), “**Inflation forecast targeting: Implementing and monitoring inflation targets**”, *European Economic Review*, No 41(6), pp. 1111-1146.

Svensson, Lars E. O. (1999), “**Inflation targeting: Some extensions**”, *Scandinavian Journal of Economics*, No 101 (3), pp. 337-361.



Svensson, Lars E. O. (2017), "**Cost-benefit analysis of leaning against the wind**", *Journal of Monetary Economics*, No 90(C), pp. 193-213.

Svensson, Lars E.O. and Woodford, Michael (2003), "**Indicator variables for optimal policy**," *Journal of Monetary Economics*, No 50(3), pp. 691-720.

Tucker, Paul M. W. (2015), "**Microprudential versus macroprudential supervision**", *remarks at the 59th Annual Conference of the Federal Reserve Bank of Boston on "Macroprudential Policy,"* Boston, 2-3 October.



Imprint and acknowledgements

Stephen G. Cecchetti is Professor and Rosen Family Chair in International Finance at the Brandeis International Business School, Research Associate at the NBER, Research Fellow at the CEPR, and Vice Chair of the ESRB Advisory Scientific Committee; Javier Suarez is Professor of Finance at CEMFI, Research Fellow at the CEPR, Research Associate at the ECGI, and Chair of the ESRB Advisory Scientific Committee. We thank all our ESRB Advisory Scientific Committee colleagues for detailed comments and suggestions; as well as the broader ESRB community, including Michal Dvořák, Stephan Fahr, Philipp Hartmann, Tuomas Peltonen, and especially Antonio Sánchez Serrano, for their assistance and support in the preparation of this report. All views expressed here are those of the authors.

Stephen G. Cecchetti

Brandeis International Business School; e-mail: cecchetti@brandeis.edu

Javier Suarez

CEMFI; e-mail: suarez@cemfi.es

© European Systemic Risk Board, 2021

Postal address 60640 Frankfurt am Main, Germany
Telephone +49 69 1344 0
Website www.esrb.europa.eu

All rights reserved. Reproduction for educational and non-commercial purposes is permitted provided that the source is acknowledged.

The cut-off date for the data included in this report was 31 May 2021.

For specific terminology please refer to the [ESRB glossary](#) (available in English only).

PDF ISBN 978-92-9472-208-9, ISSN 2600-2418, doi:10.2849/977783, DT-AI-21-002-EN-N