DISCUSSION PAPER SERIES

DP17721

USING GROWTH-AT-RISK TO ASSESS THE STANCE OF MACROPRUDENTIAL POLICY

Stephen Cecchetti and Javier Suarez

MONETARY ECONOMICS AND FLUCTUATIONS AND BANKING AND CORPORATE FINANCE

CEPR

ISSN 0265-8003

USING GROWTH-AT-RISK TO ASSESS THE STANCE OF MACROPRUDENTIAL POLICY

Stephen Cecchetti and Javier Suarez

Discussion Paper DP17721
Published 02 December 2022
Submitted 26 November 2022

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

USING GROWTH-AT-RISK TO ASSESS THE STANCE OF MACROPRUDENTIAL POLICY

Abstract

The ultimate objective of macroprudential policy is to minimise the frequency and severity of economic losses arising from severe financial distress. Critically, any policymaker seeking to achieve this goal must be able to assess the stringency or laxity of their policy settings. Applying the lessons of monetary policy, we note that doing this requires a normative measure of policy stance based on a framework comprising an objective, tools, and a model linking the two. In this paper, we propose using growth-at-risk to define an overarching objective for macroprudential policy and then show how it can form the basis for a measure of policy stance. We proceed to discuss the challenges inherent in the implementation of our framework. These include the availability of appropriate data, choices among multiple tools, and gauging the uncertain responses of the economy and financial system to macroprudential policy actions.

JEL Classification: E44, E60, G01

Keywords:

Stephen Cecchetti - cecchetti@brandeis.edu
Brandeis University and CEPR

Javier Suarez - suarez@cemfi.es
CEMFI

Acknowledgements
This paper is based on a Report of the Advisory Scientific Committee (ASC) of the European Systemic Risk Board (ESRB) written by the authors (see Cecchetti and Suarez, 2021). We thank all our ASC colleagues for detailed comments and suggestions; as well as the broader ESRB community, including Michal Dvořák, Stephan Fahr, Philipp Hartmann, Tuomas Peltonen, and Antonio Sánchez Serrano. All views expressed here are those of the authors.

Using growth-at-risk to assess the stance of macroprudential policy [1]

Stephen G. Cecchetti
Brandeis International Business School and CEPR

Javier Suarez
CEMFI and CEPR

November 2022

Abstract

The ultimate objective of macroprudential policy is to minimise the frequency and severity of economic losses arising from severe financial distress. Critically, any policymaker seeking to achieve this goal must be able to assess the stringency or laxity of their policy settings. Applying the lessons of monetary policy, we note that doing this requires a normative measure of policy stance based on a framework comprising an objective, tools, and a model linking the two. In this paper, we propose using growth-at-risk to define an overarching objective for macroprudential policy and then show how it can form the basis for a measure of policy stance. We proceed to discuss the challenges inherent in the implementation of our framework. These include the availability of appropriate data, choices among multiple tools, and gauging the uncertain responses of the economy and financial system to macroprudential policy actions.

JEL Classification: E44, E60, G01

Keywords: financial stability, macroprudential policy, growth-at-risk

[1] This paper is based on a Report of the Advisory Scientific Committee (ASC) of the European Systemic Risk Board (ESRB) written by the authors (see Cecchetti and Suarez, 2021). We thank all our ASC colleagues for detailed comments and suggestions; as well as the broader ESRB community, including Michal Dvořák, Stephan Fahr, Philipp Hartmann, Tuomas Peltonen, and Antonio Sánchez Serrano. All views expressed here are those of the authors. E-mails: cecchetti@brandeis.edu and suarez@cemfi.es.

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.

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.[2] 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.[3] 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, often in the form of a board that coordinates responsibilities and policy tools across a suite of regulatory and supervisory authorities.[4] Macroprudential policies have been in place under this name only since the financial crisis of 2007-9. Partly because this is such a recent enterprise,

[2] 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).

[3] 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.

[4] In the European Union, the coordinating institution is the ESRB.

there is an active debate over how to formulate objectives, how to use the available tools, and how to structure governance – an especially delicate matter due to the diversity of agencies and tools involved. 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.

Applying some of the lessons learned from the development of the agreed-upon monetary policy framework, in this paper we discuss the challenges associated with the development of a measure of macroprudential policy stance and propose an approach based on a metric connected to an explicit conceptual framework. We provide an alternative to the current predominantly narrative approach, offering some examples as well as a perspective on how to measure the macroprudential policy stance in a more compact and systematic manner.

As an example of what is currently feasible, 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 as a proxy for financial stability and to elaborate on how to build a notion of macroprudential policy stance around such a concept and the empirical techniques available for its implementation. 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. The approach might indeed be helpful in further improving the difficult task of coordinating and assessing the cumulative effects of policies that, as in the case of the European Union, are commonly decentralised (at least in their implementation) across multiple agencies.

This paper 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, explaining why growth-at-risk provides a useful metrics in financial stability context. In Section 4 we present a simple formal model that, relying on the growth-at-risk approach, 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 Section 5 we consider implementation issues. Section 6 concludes.

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' exposures 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.

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, the model implies a 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.[5]

[5] See Svensson (1999) or Woodford (2003) for explicit formulations of the monetary policy design problem.

Figure 1
A generic macroeconomic framework

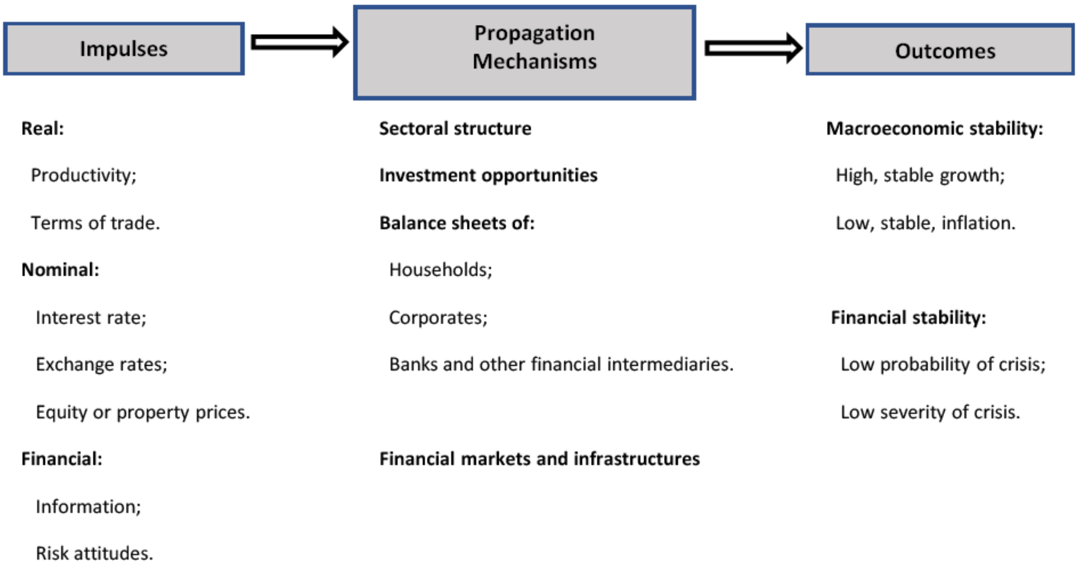

From the perspective of the generic framework, we generally cast the central bankers' problem as one where they work to meet their stabilisation objective by *reacting* to shocks which, if they were allowed to propagate, would destabilise the system. In other words, monetary policy interventions short-circuit,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

The typical 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.[6]

Turning to macroprudential policy, following Tucker (2015) we can frame the role of financial stability policymakers 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.

6 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 to us describe policies as optimally neutral, accommodative or restrictive, as well as whether they are insufficiently or excessively accommodative or restrictive.

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

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 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, in the case of banks, changing the level of capital requirements, imposing maximum loan-to-value ratios for residential

---

[8] See Hanson, Kashyap and Stein (2011) for a detailed discussion of the externalities that are the basis for macroprudential 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."

mortgages, modifying sectoral risk weights in capital requirements, and defining alternative stress test scenarios, to mention just a few. For non-bank financial intermediaries macroprudential tools are less developed but also include or might include stress tests, add-ons to liquidity requirements, and measures aimed to mitigate phenomena such as the destabilising effect of abrupt redemptions in the asset management sector or the procyclical effects of margining practices by central clearing platforms. To achieve 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 implies 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 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 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. Macroprudential 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 in this case 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 entities or the transactions in specific markets, and then fashions dedicated tools to address these risks. For example, bank regulators and supervisors use capital requirements to mitigate banks' solvency risk and loan-service-to-income limits to contain residential real estate risk, while securities markets regulators may demand that asset managers accumulate liquidity buffers to avoid spillovers arising from the fire sale of less liquid assets when facing abnormally high redemptions. 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, the current system accommodates the dispersion of the governance of macroprudential tools across authorities that exists in many jurisdictions.

Our aim is to explore the possibility of complementing this fragmented methodology with one that relies on 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.

[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."

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 remainder of this paper 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 adequately 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 many (or even most) macroprudential instruments from one day to the next. This likely delays and prolongs the impact of macroprudential 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 for banks illustrates what we mean. 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] In addition to this minimum, 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 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, existing evidence suggests that growth exhibits pronounced negative skewness and that systemic financial distress contributes

[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 (2021).

to explain the frequency and severity of adverse growth outcomes. Figure 2, taken from Cecchetti and Suarez (2021), plots the distribution of normalised average three-year growth in a large cross section of countries for two samples: the first covering years 1870-2017 (long sample) and the second covering years 1960-2017 (short sample).[12] 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 banking 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.[13]

Figure 2

Distribution of normalised average three-year growth

(percentages)

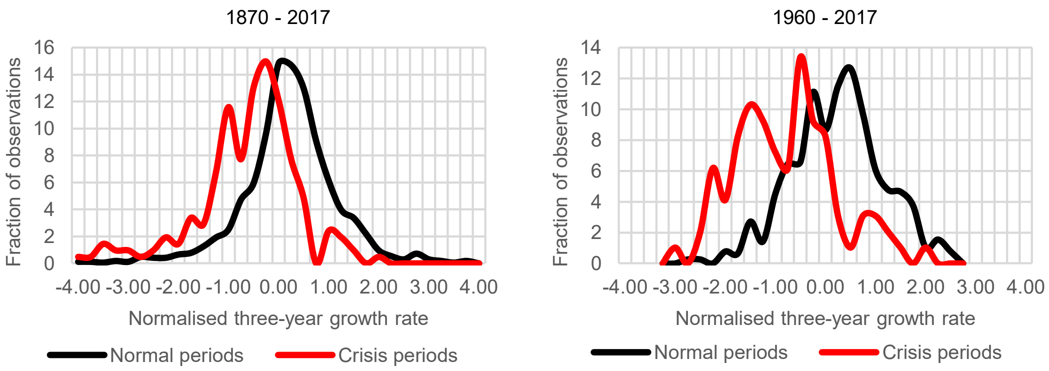

Sources: Maddison Project Database (2020); Baron, Verner and Xiong (2020; and authors' calculations. Notes: Data are deviations from the country mean of non-overlapping three-year average growth rates in standard deviation units. Countries are Argentina, Australia, Austria, Belgium, Brazil, Canada, Chile, 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.

[12] Growth rates are computed from the output per capita of the Maddison Project Database and the dating of banking crises is taken from Baron, Verner and Xiong's (2020) recently published chronology. The full dataset covers 46 countries from 1870 to 2018 and includes 207 crisis episodes. 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).

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

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. These findings are consistent with Laeven and Valencia (2018) who 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.

To connect the patterns found in the data to the growth-at-risk approach, consider the stylised distribution of output growth shown in Figure 3. 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 includes 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.[14] For future reference, we also define *growth-given-stress,* $y^S(q)$, as the expected growth rate conditional on being below the threshold $y^R(q)$.

Figure 3
Stylised probability density of output growth

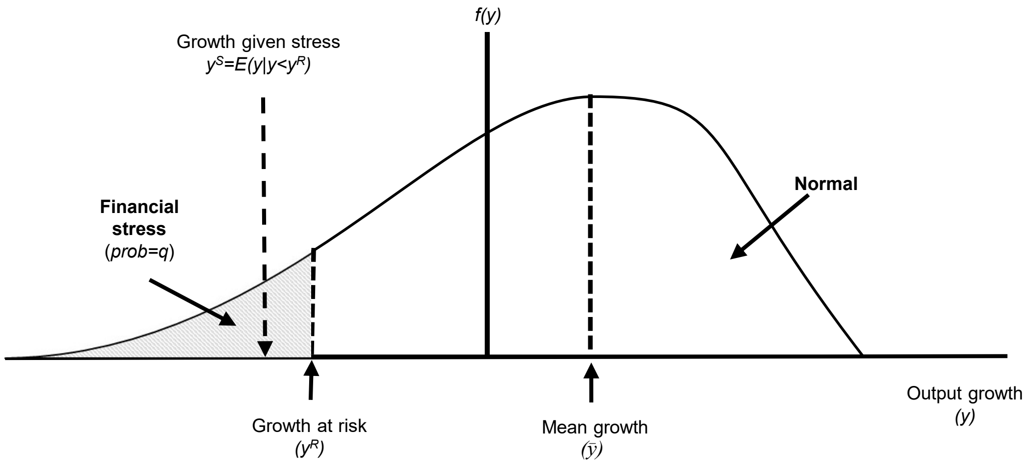

[14] See Wang and Yao (2001), Cecchetti (2008), and Adrian, Boyarchenko and Giannone (2019) for seminal applications of the growth-at-risk concept.

We note that for a reasonable 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.[15]

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.

A framework that relies on either growth-at-risk or growth-given-stress as proxies for financial stability has the 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 might happen, note that standard empirical analyses in other policy fields, including monetary policy, 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).[16] This implies 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.

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

[16] Recent applications of the growth-at-risk approach and related approaches include Caldera Sánchez and Röhn (2016), De Nicolo and Lucchetta (2017), Duprey and Ueberfeldt (2018), Falconio and Manganeli (2020), Gadea Rivas, Laeven and Perez-Quiros (2020), and Galán (2020).

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 in which 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 (2022) 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 A provides more details.[17]

[17] Suarez (2022) 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 A
Optimal policy in the CARA/normal case

Suarez (2022) 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

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

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)]

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

and

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

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 (A.1), subject to (A.2) and (A.3), yields a rule in which policy is a linear function of systemic risk:

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

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

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

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.

We note two points. 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 shows, it is straightforward 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 (2022), Appendix A.1, 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 $\rho = 4$, $\omega = 3.12$.

(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 (A.2) and (A.3), as well as the macroprudential policy reaction function (A.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 5.2 of Suarez (2022) 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.

[END OF THE BOX]

At this stage it is worth taking a moment 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 (2022). 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 their ability to reduce financial stability risks without sacrificing growth 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 A as an approximation to potentially non-linear relationships in the range over which policy entails a trade-off.

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 refers to as "forecast 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.

The Suarez (2022) 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 A). 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 A.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 A.5) from 2 to 6. The horizontal axis shows the relative impact of policy, while the vertical axis is the optimal target distance.

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

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

Figure 4

Optimal target distance from mean growth ($\bar{y}$) to growth-at-risk (y^RA. Threshold probability of stress q = 10%

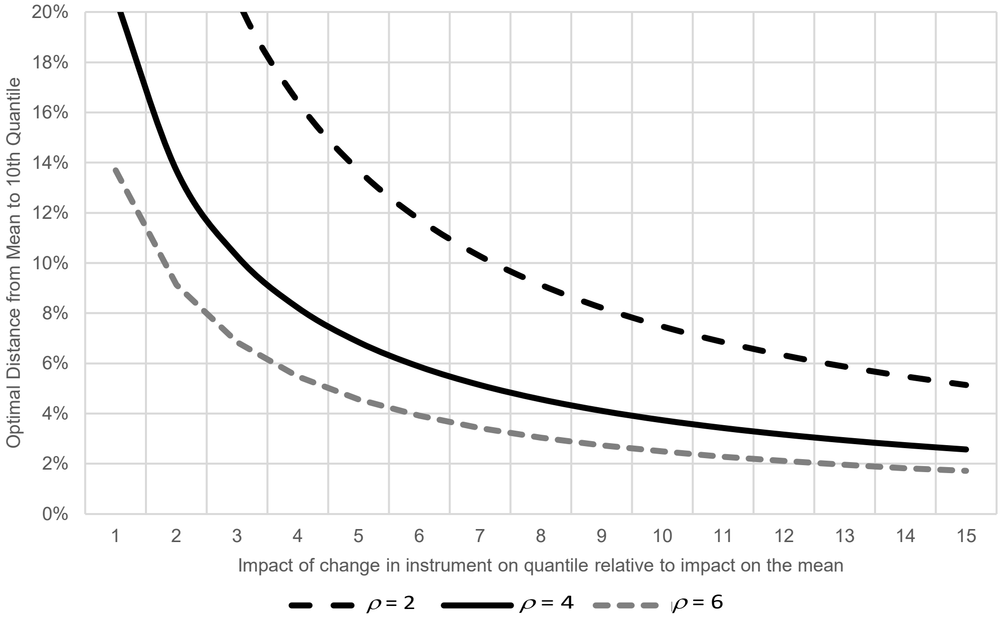

B. Relative risk aversion $\rho = 4$

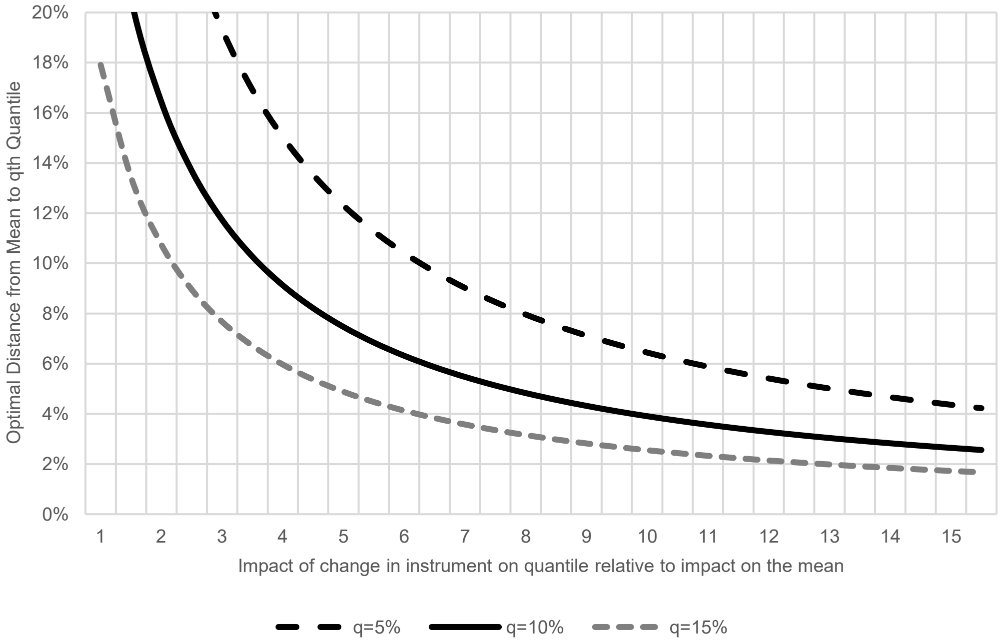

Sources: Authors' calculations based on equation (A.5) in Box A.

When $\rho = 4$ and the relative impact (γ_q/γ) equals 5 (a value roughly consistent with the results reported in Galán (2020)), the optimal target distance $(\bar{y} - y^R)^*$ is 6.84 percentage points. This number rises as risk aversion declines. When $\rho = 2$ and the relative impact remains at 5, the optimal target distance rises to 13.69 percentage points. In the bottom panel of Figure 4 we set 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.

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.

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. Until July 2021, the ECB stated its objective as price stability, which was 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. 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 describe earlier in this paper, 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; and 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.

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, for banks, increases in the countercyclical capital buffer (CCyB) 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 only makes sense for the objective to be at a longer 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. In securities markets, some polices might have a more immediate impact (e.g., temporary exemptions to clearing duties or changes in rules regarding the acceptability of assets as collateral by central counterparties) but others will similarly affect the system only over time (e.g., modifying underwriting standards in debt markets).

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.[20] 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) data 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 those 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 might provide an argument for preferring the 10th to the 5th quantile (and relative 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 is better, growth-at-risk or growth-given-stress? From a conceptual perspective the latter might have the advantage of taking the full form of the lower tail of the growth distribution into account and not just the point that corresponds to the reference low quantile. Depending on the shape of the distribution at its tail, a fixed growth-at-risk is compatible with many different values of the growth-given-stress, i.e., the growth conditional on the system being under stress. However, focusing on growth-at-risk can

[20] 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). Section 5.4 in Suarez (2022) provides a simplified treatment of this issue.

be preferable from a practical empirical perspective. Computing growth-given-stress requires estimating the area under the entire lower tail, and the absence of data to pin down the density at very low quantiles makes 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 plays a role in the choice of the horizon itself. Second, we have more experience with some tools than with others. For example, in the banking sector, changes in maximum loan-to-value ratios for residential mortgages are more common than adjustments to the CCyB or changes in bank asset concentration limits. If a tool shows no variation, then available data will be silent on its effectiveness. Third, there is a possibility that the assumed policy trade-off may not apply to all settings of each policy tool (e.g., because some tools have a negative impact on mean growth at high levels of activation but not at low ones). 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.[21]

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.

Finally, 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

[21] 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.

determine which are substitutes and which are complements, so that we can employ such tools in the best possible combinations, equating their marginal effectiveness.[22] Second, as always, policymakers need to avoid reacting to "noise". Given how underdeveloped data systems are for some parts of the financial system (especially for 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.[23] Third, as should be clear from our discussion, the policy target is likely to differ across jurisdictions. Attitudes toward risk (or society's aversion to financial instability) will diverge, as will the structure of financial systems 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.).

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 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 paper 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 concluding, it is important that we provide a few warnings. First and foremost, the purpose of this paper is to provide a perspective on the problems faced by macroprudential policymakers – including those related to the existence of several agencies involved in the pursuing of multiple intermediate objectives and the management of many tools. 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 survive in more complex, more detailed, and more

[22] See Suarez (2022), Section 4.2 for a general discussion of this problem.

[23] 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.

realistic 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, attitudes toward risk taking change 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 paper would advise against such policy actions.

To conclude, the developments summarized in this paper constitute the beginning of 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), 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, Hacioglu 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.

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.*

Caldera Sánchez, Aida and Röhn, Oliver (2016), "**How do policies influence GDP tail risks?**", *OECD Economics Department Working Papers*, No 1339.

Cecchetti, Stephen G. (2008), "**Measuring the macroeconomic risks posed by asset price booms**", in Campbell, J.Y. (ed.), *Asset Prices and Monetary Policy, Chicago, Ill*, University of Chicago Press for NBER, 9-43.

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, 89-135.

Cecchetti, Stephen G. and Suarez, Javier (2021), "**On the stance of macroprudential policy**", ESRB, Reports of the Advisory Scientific Committee, No 11, December.

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.

De Nicolo, Gianni and Lucchetta, Marcella (2017), "**Forecasting tail risks**", *Journal of Applied Econometrics,* No 32, 159-170.

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

Duprey, Thibaut and Ueberfeldt, Alexander (2018), "**How to manage macroeconomic and Financial stability risks: A new framework**", *Bank of Canada Staff Analytical Note*, No 11.

Elenev, Vadim, Landvoigt, Tim and Van Nieuwerburgh, Stijn (2021), "**A macroeconomic model with financially constrained producers and intermediaries**", *Econometrica*, No 89, 1361-1418.

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.

Gadea Rivas, Maria Dolores, Laeven, Luc and Pérez-Quirós, Gabriel (2020), "**Growth-and-risk trade-off**", *CEPR Discussion Papers*, No 14492.

Galán, Jorge E. (2020), "**The benefits are at the tail: uncovering the impact of macroprudential policy on GDP growth**", *Journal of Financial Stability, No 58, 100831*.

Hanson, Samuel G., Kashyap, Anil K. and Stein, Jeremy C. (2011), "**A macroprudential approach to financial regulation**", *Journal of Economic Perspectives*, No 25(1), 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), 259-278.

King, Mervyn (1999), "**Challenges for monetary policy: new and old**", Bank of England, speech given on 27 August.

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.

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.

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.

Suarez, Javier (2022), "**Growth-at-risk and macroprudential policy design**", *Journal of Financial Stability, No 60, 101008*.

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. 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.

Wang, Yan and Yao, Yudong (2001), "**Measuring economic downside risk and severity: Growth at risk**", *World Bank Policy Research Working Papers*, No 2674.