Growth-at-risk and macroprudential policy design

Javier Suarez

CEMFI and ESRB-ASC

January 2021

Abstract

This paper explores a potential application of the empirical growth-at-risk (GaR) approach to the design and assessment of macroprudential policies. It considers a simple linear specification of the empirical GaR approach in combination with a linear-quadratic social welfare criterion that rewards expected GDP growth and penalizes the gap between expected GDP growth and GaR. Akin to the mean-variance approach in portfolio theory, if the growth rate follows a normal distribution, such welfare criterion can be microfounded as consistent with expected utility maximization under preferences for GDP levels exhibiting constant absolute risk aversion. Under the baseline formulation, the optimal policy rule is linear in the risk indicator, with a sensitivity to the risk indicator which is independent of the risk preferences embedded in the welfare criterion. Such sensitivity depends directly on the impact of risk on the gap between expected growth and GaR, and inversely on the effectiveness of policy in reducing such gap. The optimal macroprudential policy targets a gap between expected growth and GaR which does not depend on the time-varying risk indicator but on the cost-effectiveness of macroprudential policy and the risk preference parameter. Implications for the use of the empirical GaR approach in developing a metrics for macroprudential policy stance as well as multiple extensions are discussed.

JEL Classification: G01, G20, G28

Keywords: macroprudential policy, policy stance, growth-at-risk, quantile regressions

1   Introduction

This paper is motivated by the growing attention received by the concept of growth-at-risk in the assessment of macroprudential policies. The concept was born as a natural extension to the assessment of systemic risk of value-at-risk, a popular risk management concept. In risk management, the value-at-risk of a given portfolio position is the critical level of the estimated distribution of the possible losses over a reference horizon that realized losses will not exceed with a high probability such as 95% or 99% known as the confidence level of the assessment. Estimating the value-at-risk allows the portfolio holder to assess, for example, which capital position would be needed for the absorption of the potential losses over the reference horizon with such confidence level of probability. From a statistical viewpoint, the value-at-risk of a portfolio is just the estimate of a low quantile (5% or 1% in the above examples) of the distribution of the value of the portfolio by the end of the reference horizon.

Parallel to the concept of value-at-risk, the growth-at-risk of an economy over a given horizon is a low quantile of the distribution of the (projected) GDP growth rate over such horizon. That is, the growth rate such that the probability that the realized growth rate falls below it equals a low benchmark level such as 15%, 10% or 5%.[1] Opposite to the standard macroeconomic focus on the expected value (and, perhaps, the variance) of aggregate output growth, looking at low quantiles of such growth implies, as in risk management, caring about the severity of potential adverse outcomes. Additionally to measuring such severity, the approach can provide information on the variables that determine the probability or severity of bad outcomes, including policy variables that might then be used to influence or "manage" such aggregate risk.

The rising popularity of growth-at-risk in financial stability and macroprudential policy assessments is driven by demand and supply factors. From the demand side, macroprudential policy assessment and design is in bad need for a quantitative framework that provides a baseline for policymakers' discussion, decision-making, and communication with the public similar to those provided by standard macroeconomic models, targets, and indicators in the fields of monetary policy or fiscal policy. For macroprudential policies, the multiplicity of tools, the multidimensional nature (and still vaguely defined concept) of systemic risk, data limitations, and the relatively short historical experience with the use of most policy tools pose significant challenges for the development of such framework. As a result, macroprudential policy is largely assessed and developed following a piece-meal approach (that is,

---

[1] The use of lower implied confidence levels (85%, 90%, 95% in the examples above) in growth-at-risk than in value-at-risk is partly related to the fact that GDP is not observed at frequencies that allow for a precise estimation of extremely low quantiles.

splitting the task by sector, by tool, by risk or detected vulnerability, or by a combination of approaches) and relying on the expert judgement for the qualitative integration of the pieces. While the aim is to cover the whole financial system, the resulting assessment is often less complete, integrated, systematic, and quantitative than in other policy fields.

From the supply side, the impulse to the use of the concepts of GDP-at-risk (Cecchetti, 2008) and growth-at-risk (Adrian et al., 2018) is mainly empirical. It is related to the availability of econometric techniques that extend regression analysis (single dependent variable models, panel data models, vector auto-regressive models) to quantiles and their use in macroprudential applications by a growing number of authors. Quantile regression techniques allow to shift the focus from modeling the conditional mean of the dependent variable to modeling the conditional quantiles (and thereby the whole conditional distribution) of the dependent variables.

Quantile regressions allowed Cecchetti and Li (2008) to use the concept of GDP-at-risk as an empirically-viable summary measure of the impact of asset price booms on financial stability. This approach was further developed and promoted by the influential paper of Adrian, Boyarchenko, and Giannone (2019), which shows that the lower quantiles of the distribution of the US GDP growth rate fluctuate more and are more influenced by financial conditions than the upper quantiles, thus supporting the focus of macroprudential surveillance and policies on such lower quantiles. Adrian et al. (2018), which coined the acronym GaR for growth-at-risk (henceforth also used in this paper), documented the "term structure" of GaR and suggested the existence of an intertemporal trade-off whereby some policies might improve GaR at medium and long horizons but at the cost of damaging GaR (or expected growth) at shorter horizons. Other contributions following a quantile-regression approach to the analysis of growth vulnerabilities and their relationship with financial conditions and macroprudential policies include Caldera-Sánchez and Röhn (2016), De Nicolo and Lucchetta (2017), Duprey and Ueberfeldt (2018), Aikman et al. (2019), Prasad et al. (2019), Galán (2021), and Chavleishvili et al. (2020). The empirical approach to GaR has also been embraced in part of its ongoing work by the Expert Group "Macroprudential Stance - Phase II" of the European Systemic Risk Board (ESRB).

Relative to other indicators of financial stability, GaR features the advantage of having an explicit and intuitive statistical interpretation and being measured in the same units as GDP growth, the most universal summary indicator of an economy's overall performance. Hence, quantitative contributions around the concept of GaR are followed with great interest (and some skepticism too) by the institutions involved in the assessment and design of macroprudential policies. Many see in them a promising step in the development of the

missing integrated quantitative framework mentioned above.[2] However, as further discussed in Cecchetti and Suarez (2021), existing empirical efforts still lack a clear fit into an explicit policy design problem of the type considered for other macroeconomic policies (e.g., in the derivation of an optimal monetary policy rule).

This paper aims to fill this gap by digging into the potential application of the empirical GaR approach to the design and assessment of macroprudential policies. Relying on a stylized representation of the type of equations that the quantile regression approach may deliver, the paper studies how macroprudential policy could be designed and evaluated using a linear-quadratic social welfare criterion that rewards expected GDP growth and penalizes the gap between expected GDP growth and GaR. It is shown that, in specific environments, such welfare criterion can be microfounded as consistent with expected utility maximization under risk averse preferences for GDP levels. The paper characterizes the properties of the optimal macroprudential policy rules in the basic setup and a number of relevant extensions. Implications are drawn on the possibility of assessing macroprudential policy stance with a metric emanated from the estimated equations of the empirical GaR approach.

Under the baseline formulation, the optimal policy rule is linear in the risk indicator, with a sensitivity to the risk indicator which is independent of the risk preferences embedded in the welfare criterion. Such sensitivity depends directly on the impact of risk on the gap between expected growth and GaR, and inversely on the effectiveness of policy in reducing such gap. The optimal macroprudential policy targets a gap between expected growth and GaR which does not depend on the time-varying risk indicator but on the cost-effectiveness of macroprudential policy and the risk preference parameter.

The explored variations of the basic setup cover cases with non-linearities in the impacts of policy variables and risk variables on the relevant outcomes, multiple policy variables, and discrete policy variables. An important extension shows the compatibility of the GaR framework (and the main insights from the basic formulation) with the view that macroprudential policy involves various well-identified intermediate objectives each of which can be associated with one or a subset of targeted policy tools. Additional discussions deal with policies that seem to involve no trade-off between mean growth and GaR, the treatment of country heterogeneity, the interaction with other policies, and the possibility of reformulating the

---

[2]Skeptics have doubts about the feasibilty and/or desirability of such an integrated approach, since they think that the multidimensionality of macroprudential policy cannot be subsumed by looking at a single aggregate indicator such as GaR. Instead, a policymaker in this field might have to keep track of a welfare criterion that directly combines (intermediate) objectives along the many dimensions of systemic risk and takes into account how (potentially interacting) policies affect all such (intermediate) objectives. In the EU, Recommendation ESRB/2013/1 establishes five intermediate objectives for macroprudential policy.

analysis around the concept of growth-given-stress rather than GaR.

The paper is structured as follows. Section 2 provides a basic linear formulation of the empirical GaR approach. Section 3 develops the welfare criterion used for optimal policy design in such setup, and derives and establishes the properties of the optimal macroprudential policy rule. Section 4 discusses the implications of the results for the assessment of macroprudential policy stance (that is, how the estimates associated with empirical GaR regressions could help inform about the stance of macroprudential policy). Section 5 develops several extensions of the basic setup, generalizing its results to a variety of empirically relevant cases as well as to the case in which macroprudential policy comprises several intermediate objectives which can be addressed with targeted tools. Section 6 contains some further discussion of the proposed analytical framework and the results. Section 7 concludes the paper. The Appendix contains the microfoundations of the GaR-based welfare criterion used for the design of the optimal policies and discusses the extent to which, when departing from normality, focusing on the low tail of the GDP growth distribution over a given horizon might have advantages over an alternative focus on just the conditional mean and conditional variance of such growth distribution.

2   A basic formulation of the empirical GaR approach

A quantile regression approach can deliver equations for arbitrary quantiles of GDP growth over relevant horizons. Consider a stylized representation of such approach that consists of two estimated equations: one for the mean (or perhaps the median) of GDP growth, denoted $\bar{y}$, and another one for a relevant low quantile of GDP growth, y_c, where c is the so-called *confidence level* at which GaR is measured. By definition, y_c satisfies

$$\Pr(y \leq y_c) = c. \tag{1}$$

which means that the probability of experiencing growth rates lower than y_c over the relevant horizon is just c. The confidence level c can be thought to be 5% or 10% so that y_c reflects how bad growth may be under adverse circumstances typically associated with systemic distress.

For the purposes of illustration, consider the simple case in which the quantile regression approach delivers conditional forecast equations for *GaR* y_c and *expected growth* $\bar{y}$ of the form

$$y_c = \alpha_c + \beta_c x + \gamma_c z, \tag{2}$$

and

$$\bar{y} = \alpha + \beta x + \gamma z, \tag{3}$$

where x is a unidimensional measure of *risk* (e.g., a measure of excessive credit growth or any other index reflecting the accumulation of financial imbalances) and z is a unidimensional macroprudential *policy* variable (e.g., a bank capital-based measure such as the countercyclical capital buffer –CCyB– of Basel III).[3] Assume that the endogeneity of z has been treated well enough to allow for γ_c and γ to be interpreted as the causal impact of variations in z on GaR and expected growth, respectively.

Assume further that

$$\beta_c < \min\{0, \beta\} \text{ and } \gamma < 0 < \gamma_c. \tag{4}$$

In words, the risk indicator x has a negative impact on GaR and a less negative (or even positive) impact on expected growth, while the policy variable z has a positive impact on GaR but a negative impact on expected GDP growth.[4] These last properties imply that the policy measured by z involves a trade-off.[5]

For example, if x measures excessive credit growth and z is the CCyB rate, the trade-off would arise if increasing the CCyB rate can reduce the risk associated with, say, a credit boom (e.g., the probability and implications of an abrupt reversal) but, at the same time, has a negative contractive impact on aggregate demand and, hence, on the central outlook.

Finally, assume that the ranges of variation of y and z together with the values of the intercepts α and α_c guarantee $y_c < \bar{y}$ over the relevant range (otherwise, the linearity in (2) and (3) might lead to $\bar{y} < y_c$ which would not make sense for low values of c).

3   Social preferences and the optimal policy rule

To further illustrate the policy trade-offs derived from the empirical GaR formulation, suppose the policy maker has preferences than can be represented by the social welfare function

$$W = \bar{y} - \frac{1}{2}w(\bar{y} - y_c)^2, \tag{5}$$

---

[3]An advance reader might easily extend some of the derivations and claims contained in this note to the cases in which x and z are vectors of risk indicators and policy variables, respectively. See Sections 5 for extensions of the basic formulation that deal with multiple policy variables. Section 6.3 considers the interaction with policies other than macroprudential ones.

[4]The linear specification implies that z affects monotonically y and y_c. What really matters for the validity of the analysis below is that this is locally true over the relevant range of variation in z. Otherwise the specification could be modified by redefining z as a suitable non-monotonic transformation of the policy variable.

[5]See Adrian et al. (2018) or Galán (2021) for cross-country evidence of such trade-off.

where $w > 0$ measures the aversion for financial instability, which here is proxied by the magnitude of the quadratic deviations of GaR with respect to expected growth.

As shown in Section A.1 of the Appendix, in the particular case in which GDP growth follows a normal distribution, the welfare criterion in (5) can be justified as consistent with the maximization of the expected utility of a representative risk-averse agent whose utility depends on GDP levels. Specifically, if the agent has preferences for GDP levels exhibiting constant absolute risk aversion (CARA), say λ, then (5) provides an exact representation of such preferences under a value of w which is directly proportional to λ.

Of course, if y is normally distributed, social preferences and the policy problem could have also been formulated in the usual mean-variance terms of portfolio theory, with an equation describing the dependence of the standard deviation of the growth rate σ_y on x and z replacing (2) (see Section A.2 of the Appendix for details). What this means is that the true advantages of adopting a GaR approach (instead of a mean-variance approach) in the formulation of the macroprudential policy problem must come from the fact that, in reality, (i) the conditional distribution of y is not Gaussian, and (ii) as documented in recent empirical work, the financial factors and policy tools on which macroprudential policy focuses affect the conditional low quantiles of the true growth distribution in a stronger and better identifiable manner than its conditional variance. From this perspective, an advantage of the quantile-regression approach to the modelling of the quantile y_c is that it does not require assuming a specific distribution for the conditional quantile. That is, nothing prevents the estimated version of equations (2) and (3) to capture features such as the potential left skewness of the true conditional distribution of the GDP growth rate.

Beyond the exact expected-utility microfoundations of the specific normal case, the welfare criterion in (5) could be defended also in heuristic or axiomatic terms as the representation of the preferences of a policy maker that faces a trade-off between improving mean outcomes and reducing the severity very bad outcomes. An interesting feature of (5) from such perspective is that the dislike for "very bad outcomes" is proportional to the square of the distance between the bad outcomes y_c and the mean outcomes $\bar{y}$, where the latter would play the role of a reference level (or status quo point) similar to those emphasized in some non-expected-utility formulations of agents' preferences for risk. Specifically, from the perspective of prospect theory (Kahneman and Tversky, 1979), the coefficient w in (5) could be interpreted as capturing loss aversion rather than risk aversion.[6]

---

[6] The asymmetric focus on low tail losses can also be related to Fishburn (1977) that explores preferences in which the decision maker is averse to obtaining below-target payoffs. Kilian and Manganelli (2008) analyze the decision problem of a central banker using that approach. In a related vein, Svensson (2003) considers a monetary policy problem under preferences that asymmetrically penalize extreme events.

3.1 The optimal policy rule

An optimal macroprudential policy conditional on x would thus maximize W given x. That is, it would be characterized by the policy rule

$$z(x) = \arg\max_z W(x, z), \tag{6}$$

where $W(x, z)$ describes W as a function of the risk indicator x and the policy variable z after taking into account (2) and (3).

If the optimal policy is interior, it must solve the following first order condition (FOC):

$$\frac{\partial \bar{y}}{\partial z} - w(\bar{y} - y_c) \left(\frac{\partial \bar{y}}{\partial z} - \frac{\partial y_c}{\partial z} \right) = 0, \tag{7}$$

which uses the chain rule in (5). From (2) and (3), this FOC can be written as

$$\gamma - w(\alpha + \beta x + \gamma z - \alpha_c - \beta_c x - \gamma_c z)(\gamma - \gamma_c) = 0. \tag{8}$$

Solving for z leads to the *macroprudential policy rule*

$$z(x) = \phi_0 + \phi_1 x, \tag{9}$$

with

$$\phi_0 = \frac{\alpha - \alpha_c}{\gamma_c - \gamma} + \frac{\gamma}{w(\gamma_c - \gamma)^2} \tag{10}$$

and

$$\phi_1 = \frac{\beta - \beta_c}{\gamma_c - \gamma}. \tag{11}$$

Under our assumptions, the *intercept* of the policy rule ϕ_0 can in principle have any sign since it is the sum of a first term which will most typically be positive (specifically if $\alpha - \alpha_c > 0$) and a second term which is negative (since $\gamma < 0$). Yet, ϕ_0 is intuitively increasing in the policy maker's preference for financial stability w (since the absolute size of the negative term declines with w) and also increasing in the difference $\gamma_c - \gamma > 0$, which measures the *effectiveness* of the policy variable in reducing the *gap* between expected growth and GaR, $\bar{y} - y_c$.[7]

---

[7]For instance, in the polar case in which the policy maker has absolute preference for financial stability ($w \to \infty$), the intercept would become just $(\alpha - \alpha_c)/(\gamma_c - \gamma)$ and lead to a solution with $\bar{y} = y_c$ (which, although unrealistic in practice, is mathematically feasible given the linearity of (2) and (3) in x and z). In the other polar scenario with no preference for financial stability ($w \to 0$), we would have $\phi_0 \to -\infty$, implying that the policy maker would choose the lowest possible value of z, since under the linear specification of (3) this is the way to maximize expected growth (albeit at the cost of minimizing GaR).

Interestingly, the parameter ϕ_1 which measures the *responsiveness* of the optimal policy to variations in the risk indicator x is positive and independent of the preference parameter w. So in this setup, policy makers with different preferences for financial stability would differ in the level at which they use macroprudential policy but not in the extent to which they modify their policies in response to changes in the risk assessment. Such optimal policy responsiveness is directly proportional to the *impact of risk* x on the gap between expected growth and GaR ($\beta - \beta_c$, which is positive under (4)) and inversely proportional to the *effectiveness of policy* z in reducing such gap ($\gamma_c - \gamma$, which is also positive under (4)).

While the empirical GaR approach, namely estimating equations (2) and (3), does not per se allow to estimate the policy parameter w, it allows to directly estimate the optimal policy responsiveness parameter ϕ_1 and its components $\beta - \beta_c$ and $\gamma_c - \gamma$. Also, from the above reasoning, it also allows to represent the optimal policy rule for different illustrative values of the preference parameter w.

3.2 Graphical illustration

Further understanding of the interaction between the policy trade-offs involved in (2) and (3) and the preferences reflected in (5) can be obtained by depicting the frontier of pairs of y_c and $\bar{y}$ that can be reached, for a given value of the risk variable x, by varying the policy variable z. Mathematically this *conditional policy frontier* (for a given x) is defined by the line:

$$\bar{y} = \left(\alpha - \frac{\gamma}{\gamma_c}\alpha_c\right) + \left(\beta - \frac{\gamma}{\gamma_c}\beta_c\right) x + \frac{\gamma}{\gamma_c}y_c, \tag{12}$$

which is downward sloping in $(y_c, \bar{y})$ space. Figure 1 depicts the policy frontier for a given value of x. The point $(y_c(x,0), \bar{y}(x,0))$ corresponds to the case in which $z = 0$. Intuitively choosing $z > 0$ allows to reach higher values of y_c but at the cost of lowering $\bar{y}$.[8]

The preferences in (5) describe a map of indifference curves in $(y_c, \bar{y})$ space that are convex parabolas that reach their minima on the ray $y_c = \bar{y}$. The map makes economic sense to the left of such ray. Intuitively, for $w > 0$, on each indifference curve, any decline in y_c should be compensated with an increase in $\bar{y}$ so as to keep the welfare level unchanged. Moreover, for a given decline in y_c, the required compensating increase in $\bar{y}$ increases with

---

[8]Under the assumed linearity, there is nothing special about $z = 0$ but in specific applications one could normalize the policy variable so that it means something, e.g., the historical mean or "normal stance" of the corresponding policy (then $z < 0$ would represent a stance looser than normal and $z > 0$ a stance tighter than normal). For some policy instruments there may be a natural lower bound to z, e.g. a CCyB rate of Basel III cannot be negative (although in practice there are instances of capital forbearance that might be similar to having $z < 0$ for such instrument). Explicit consideration of such bounds would raise complications regarding occasionally binding constraints familiar in other contexts.

the distance from the ray $\bar{y} = y_c$. This explains why the FOC (7) includes the term $w(\bar{y} - y_c)$, which accounts for the marginal cost of financial instability.

The optimal policy $z(x)$ is the choice of z that leads to maximum welfare on the corresponding conditional policy frontier. That is, $z(x)$ is the policy level that leads to the point where the conditional policy frontier is tangent to the map of indifference curves. From the determinants of the slopes of such curves it follows that, other things equal, a policy maker with a stronger preference for financial stability will choose combinations of $(y_c, \bar{y})$ on the frontier that involve lower $\bar{y}$ and higher y_c, that is, a lower gap between expected growth and GaR.

3.3 Optimal target gap property

What happens when the risk indicator x moves? From (12), changes in the risk indicator x shift the policy frontier in parallel (just another implication of the linear formulation). In the most plausible situation in which risk does not increase expected growth too much (formally, when $\beta < \gamma\beta_c/\gamma_c$, where $\gamma\beta_c/\gamma_c$ is positive under (4)), a rise in x shifts the policy frontier down, necessarily worsening the terms of choice for the policy maker. In the alternative

scenario with $\beta > \gamma\beta_c/\gamma_c$, risk has such a strong effect on expected growth that it moves the policy frontier up.

In both cases, however, the optimal policy rule (9) implies that an increase in risk leads to a tightening of the policy decision z, indicating that the fall in y_c that would occur if policy were not adjusted is, at least partly, offset by increasing z. When risk has a positive marginal impact on expected growth ($\beta > 0$), the optimal policy response will diminish the raw positive effect of the risk indicator x on mean growth $\bar{y}$. When risk has a negative marginal impact on expected growth ($\beta < 0$), the optimal policy response will still aim to offset its even more negative effect on y_c by lowering mean growth $\bar{y}$ beyond what the raw negative effect of the risk indicator x would imply.

Mathematically, one can see the final impact of changes in the risk indicator x on y_c and $\bar{y}$ by substituting the optimal policy rule $z(x)$ into (2) and (3), which leads to

$$y_c = (\alpha_c + \beta_c\phi_0) + (\beta_c + \gamma_c\phi_1)x = (\alpha_c + \beta_c\phi_0) + \frac{\gamma_c\beta - \gamma\beta_c}{\gamma_c - \gamma}x, \tag{13}$$

and

$$\bar{y} = (\alpha + \beta\phi_0) + (\beta + \gamma\phi_1)x = (\alpha + \beta\phi_0) + \frac{\gamma_c\beta - \gamma\beta_c}{\gamma_c - \gamma}x. \tag{14}$$

Interestingly, the coefficient of the risk indicator x is identical in these two equations, which implies that the optimal policy rule would keep constant the gap between expected growth and GaR:

$$\bar{y} - y_c = (\alpha - \alpha_c) + (\beta - \beta_c)\phi_0 = \frac{1}{w}\frac{(-\gamma)}{\gamma_c - \gamma} = \frac{1}{w}\frac{1}{1 + \gamma_c/(-\gamma)} \tag{15}$$

Notice that this *target gap* is positive under (4) since $\gamma < 0$. The target gap is decreasing in the preference for financial stability w and increasing in the *marginal growth-gap rate of transformation* implied by the policy frontier, $(-\gamma)/(\gamma_c - \gamma)$, which can be rewritten as $1/[1 + \gamma_c/(-\gamma)]$ to better visualize its negative dependence with respect to the *relative cost-effectiveness* of the policy variable: the ratio of the c-quantile improvement effect $\gamma_c > 0$ to the mean reducing effect $-\gamma < \gamma_c$.

In fact, this "constant target gap" property can be directly obtained from the FOC in (7), which can be rearranged as

$$\bar{y} - y_c = -\frac{1}{w}\frac{\frac{\partial \bar{y}}{\partial z}}{\frac{\partial \bar{y}}{\partial z} - \frac{\partial y_c}{\partial z}} = \frac{1}{w}\frac{1}{1 + \gamma_c/(-\gamma)}. \tag{16}$$

Graphically this implies that changes in the risk indicator x and the optimal policy response under $z(x)$ describe a linear *expansion path* in $(y_c, \bar{y})$ with slope equal to one. So starting from the optimal policy identified in Figure 1 for a particular risk level x, changes in x will lead to reach combinations $(y_c, \bar{y})$ on the line with slope one that goes through that point.

11

More specifically, when risk does not increase expected growth too much (that is, in the case $\beta < \gamma\beta_c/\gamma_c$ already described above), the coefficient of x in the reduced-form equations (13) and (14) is negative. Thus, when the risk indicator increases and policy responds optimally, GaR and expected growth *deteriorate* by the same amount, so as to keep the gap between expected growth and GaR, $\bar{y} - y_c$, constant. This is the case depicted in Figure 1 where the policy frontier under $x' > x$ lies on the left of the one for x.

Otherwise (that is, when $\beta > \gamma\beta_c/\gamma_c$), the coefficient of x in (13) and (14) is positive so rises in x and the optimal policy response lead GaR and expected growth to *improve* by the same amount but, again, keeping $\bar{y} - y_c$ constant.

An important corollary of these findings is that, under the specified preferences, macro-prudential policy should not target a constant GaR or to make GaR be above a certain lower bound but should allow the GaR target to comove (actually one-by-one!) with the expected growth estimate. In other words, these derivations suggest that the gap $\bar{y} - y_c$ is a more useful indicator of stance than each of its components separately.

4   A framework for policy assessment?

The following is a tentative list of policy-relevant outputs this approach can deliver:

1. Estimating (2) and (3) allows to positively describe the direct impact of risk x and policy z on GaR and expected growth, as well as the involved policy trade-offs.

2. The policy trade-offs can be further illustrated using a policy frontier as in Figure 1.

   (a) If evaluated at the historical mean value of x, such frontier could be called the *mean policy frontier.* If a practical application involves a discrete x, then one could select a reference value of it to represent, say, a "normal" situation.

   (b) Under the linear specification, the conditional policy frontier is just a parallel shift of the mean (or "normal") policy frontier. The relative position of the conditional frontier relative to the historical mean (or "normal") frontier may indicate whether the economy faces a state of above-normal or below-normal risk exposure.

3. If social preferences (or the preferences of the policy maker) can be described with a mean growth versus GaR welfare criterion as in (5), then

   (a) The optimal policy responsiveness to the risk indicator can be measured by $\phi_1 = \frac{\beta - \beta_c}{\gamma_c - \gamma}$, as in (11). This measure is independent of the parameter w that describes the policy maker's preference for financial stability.

(b) If the policy maker follows the optimal policy rule, it will implicitly target a constant gap between expected growth and GaR, as in (16). The optimal policy gap will be decreasing in the preference parameter w and in the cost-effectiveness ratio of the policy tool $\gamma_c/(-\gamma)$. From such gap, and the estimates of the empirical GaR model, the implicit preference parameter could be recovered (inferred) from the condition

$$w = \frac{1}{\bar{y} - y_c} \frac{1}{1 + \gamma_c/(-\gamma)}. \tag{17}$$

(c) Conditional on a reference value of w, the optimal policy rule can be fully described using (6). Graphically, it can be described with the expansion path previously illustrated in Figure 1.

(d) Conditional on a reference value of w and an assessment of risk x, a graphical counterpart of the optimal policy choice can be described by depicting the conditional policy frontier and the point on it associated with the optimal policy (which is given by the intersection between such policy frontier and the expansion path).

(e) Conditional on an assessment of risk x, a policy stance could be deemed *inefficient* if leading to points sufficiently far away from the policy frontier. However, when the policy variable x is unidimensional (as in all derivations above), all choices of x are "efficient", so the concept of inefficiency is only useful with two or more policy variables.

(f) Conditional on the reference value of w and an assessment of risk x, a policy stance could be deemed *suboptimal* if sufficiently far away from the expansion path. This corresponds to an excessive distance between z and $z(x)$ or, in terms of outcomes, a gap $\bar{y} - y_c$ far enough from its target. So policy would be *too tight* if z were sufficiently higher than $z(x)$ and, equivalently, if the gap $\bar{y} - y_c$ were well below target. Conversely, policy would be *too loose* if z were sufficiently lower than $z(x)$ and, equivalently, if the gap $\bar{y} - y_c$ were well above target.

5   Extensions

This section considers several specific extensions of the basic formulation, showing the capacity of the framework to accommodate multiple variations and checking the robustness of the properties of optimal macroprudential policies to each of them.

5.1 Policy variable with non-linear effects

As partly anticipated in the prior subsection, a particularly relevant non-linearity in practice may be related to the diminishing effectiveness of the policy variable (or variables, if there are several) in improving the GaR. Another interesting case may emerge if the impact of the policy variable on expected growth is marginally increasing. So consider a generalized version of (2) and (3) with

$$y_c = \alpha_c + \beta_c x + \Gamma_c(z), \tag{18}$$

and

$$\bar{y} = \alpha + \beta x + \Gamma(z), \tag{19}$$

where the functions $\Gamma_c(z)$ and $\Gamma(z)$ satisfy $\Gamma' < 0 < \Gamma'_c$ and $\Gamma''_c < \Gamma'' < 0$. In this case, the FOC solved by an interior optimal policy can be written as

$$\Gamma'(z) - w[\bar{y}(x,z) - y_c(x,z)][\Gamma'(z) - \Gamma'_c(z)] = 0, \tag{20}$$

where the dependence of $\bar{y}$ and y_c on x and z has been made explicit to emphasize the type of non-linear equation that would have to be solved to find the optimal policy rule $z(x)$.

By rearranging (20), one can obtain an expression for the gap associated with the optimal policy very similar to (16):

$$\bar{y}(x,z) - y_c(x,z) = \frac{1}{w} \frac{1}{1 + \Gamma'_c(z)/(-\Gamma'(z))}. \tag{21}$$

However, in this case the target gap is not invariant to the risk indicator x. If x increases, other things equal, the left hand side of (21) increases, calling for an offsetting increase in z. But the right hand side of (21) is now increasing in z because, intuitively, the policy trade-off measured by the relative marginal cost-effectiveness of the policy $(\Gamma'_c(z)/(-\Gamma'(z)))$ worsens at higher levels of z. This implies that the optimal policy $z(x)$ in this case grows less-than-linearly with x and the optimal gap increases with x. In words, as risk deteriorates, the policy maker would accommodate the lower and lower cost-effectiveness of the policy tool by widening the targeted gap between expected growth and GaR.

5.2 Risk variable with non-linear effects

Consider now a situation in which the risk variable x has a non-linear impact on expected growth and GaR captured by the following modification of (2) and (3):

$$y_c = \alpha_c + B_c(x) + \gamma_c z, \tag{22}$$

and

$$\bar{y} = \alpha + B(x) + \gamma z, \tag{23}$$

where $B_c(x)$ and $B(x)$ are functions satisfying $B'(x) - B_c'(x) > 0$ and $B''(x) - B_c''(x) > 0$. In words, increases in x increase the gap between expected GDP and GaR at an increasing rate (e.g. by making the financial system more and more likely to reach a tipping point of full meltdown). In this case, the FOC of the welfare maximization problem becomes:

$$\gamma - w[\bar{y}(x, z) - y_c(x, z)](\gamma - \gamma_c) = 0, \tag{24}$$

that implicitly defines the policy rule $z(x)$. In this case, the FOC and, consequently, the policy rule are only non-linear because of the non-linear effect of x on $\bar{y}(x, z) - y_c(x, z)$. Rearranging it to solve for the optimal gap, one obtains

$$\bar{y}(x, z) - y_c(x, z) = \frac{1}{w} \frac{1}{1 + \gamma_c/(-\gamma)}, \tag{25}$$

whose right hand side is invariant to x, thus implying the same gap as when x had a linear impact on y_c and $\bar{y}$. Interestingly, in this case when the risk indicator increases, the left hand side increases more than proportionally to x, calling for a more than proportional increase in the policy variable too. So, in this setup, the policy response to rises in risk should be increasingly aggressive as risk increases.[9]

5.3 A vector of policy variables

Consider an extended version of (2) and (3) with M different continuous policy variables z_j with $j = 1, 2, ...M$ affecting linearly y_c and $\bar{y}$ with coefficients γ_{cj} and γ_j, respectively. Assume these coefficients satisfy $\gamma_j < 0 < \gamma_{cj}$ as in (4). Assume further that the variables are scaled so that $z_j = 0$ is the lowest bound applicable to all of them. In this linear world, as one can see by exploring the relevant first order conditions, there will generally be one variable dominating the others in the maximization of W. This most efficient or *preferred policy tool* j^* would be the one featuring the lowest value of what was called the marginal growth-gap rate of transformation in the single policy variable benchmark,

$$\frac{\frac{\partial \bar{y}}{\partial z_j}}{\frac{\partial \bar{y}}{\partial z_j} - \frac{\partial y_c}{\partial z_j}} = \frac{1}{1 + \gamma_{cj}/(-\gamma_j)} > 0, \tag{26}$$

that is, the policy tool with the best cost-effectiveness as measured by $\gamma_{cj}/(-\gamma_j)$. Intuitively, when $\gamma_{cj}/(-\gamma_j)$ is higher the same reduction in the gap between expected growth and GaR

---

[9]The opposite situation in which the optimal policy is decreasingly aggressive as x increases emerges if $B''(x) - B_c''(x) < 0$.

can be achieved at a lower cost in terms of expected growth. For the most efficient tool, the optimal value of z_{j*} would be the one satisfying the counterpart of equation (7). The associated policy rule would be the same as in (9) with ϕ_0 and ϕ_1 particularized to the preferred tool j^*. All elements in Figure 1 remain valid if the policy frontiers get also particularized to those obtained using the preferred tool.

All the other policy variables should remain at their lowest bound of zero. In terms of Figure 1, using an inferior tool would imply moving over policy "frontiers" that also go through the point $(y_c(x,0), \bar{y}(x,0))$ but with steeper slopes, confirming that such tool would only be able to increase y_c by causing larger declines in $\bar{y}$. Conditional on using a less cost-effective tool, equation (16) would imply that the target gap should be larger, thus accommodating the harder trade-off faced along the corresponding policy frontier.

5.4   Optimal policy mixes

The optimality of using non-trivial combinations of tools in macroprudential policy would only emerge under departures from linearity. For example, one could obtain optimal policies that involve using several tools at the same time if the effectiveness of each policy tool in reducing GaR is marginally decreasing, say, given by some functions $\Gamma_{cj}(z_j)$ with $\Gamma'_{cj} > 0$ and $\Gamma''_{cj} < 0$, or if there are complementarities between tools under some general quasi-concave function $\Gamma_c(z_1, z_2, ... z_M)$ that replaces the terms $\Sigma_{j=1}^M \gamma_{cj} z_j$ in the extended version of (2).

In such non-linear world, all policy variables activated at a strictly positive level at the optimum would satisfy a properly modified version of (7) and, consequently, (16) implying

$$\bar{y} - y_c = -\frac{1}{w} \frac{\frac{\partial \bar{y}}{\partial z_j}}{\frac{\partial \bar{y}}{\partial z_j} - \frac{\partial y_c}{\partial z_j}} = \frac{1}{w} \frac{1}{1 + \frac{\partial \Gamma_c}{\partial z_j}/(-\gamma_j)}. \tag{27}$$

Thus optimal policy mixes would feature *equalization of the marginal cost-effectiveness ratios,* $\frac{\partial \Gamma_c}{\partial z_j}/(-\gamma_j)$, across all the activated policy tools. The optimal gap between expected GDP growth and GaR would be decreasing in both the common cost-effectiveness ratio and the aversion to financial instability.

In the world of interactions between tools, the optimal gap may no longer be constant since the compound effectiveness of a given policy mix may depend on the intensity with which policies are activated. For example, if the a rise in the risk variable x calls for a more intensive use of two complementary policies that jointly exhibit some decreasing returns to intensity (akin to when complementary inputs are combined in a production function with decreasing returns to scale), then the optimal policy will accommodate (as in the single

16

policy variable case with decreasing marginal effectiveness discussed above) the decreasing effectiveness by tolerating a larger gap when the risk is high than when the risk is low.

5.5   Intermediate objectives and targeted policy tools

Current practice of macroprudential policy involves, to a large extent, a piece-meal approach. Authorities around the globe, as well as research in the field of macroprudential policy, often address the assessment and design of macroprudential policy by splitting it into separate silos. As in microprudential regulations, such silos are commonly determined by the nature of the underlying source of systemic risk (e.g. liquidity vs. solvency risk) or by the sector that originates, transmits or suffers the risk (e.g. banks vs. non-banks, commercial vs. residential real estate, et cetera); the resulting silos are typically associated with one or several dedicated policy tools (e.g. "capital-based tools for the banking sector," "liquidity-based tools for the investment management sector" or "borrower-based tools for residential real estate risk"). The practical attractiveness of the piece-meal approach stems from the difficulties to integrate under a common general equilibrium perspective and a common ultimate goal the multiple dimensions of systemic risk, the multiple potential factors contributing to financial stability (or the lack of it), and the multiple policy tools available to address these dimensions and factors. The purpose of this subsection is to show that such approach is not incompatible with the analytical framework and empirical efforts associated with the GaR approach. In fact, this approach can contribute to integrate, add-up or are least put under a common umbrella sectoral macroprudential policies that might, otherwise, be difficult to relate to each other when trying to obtain an overall notion of macroprudential policy stance.

Of course, the illustration below of the capability of the GaR approach to integrate prior piece-meal approaches to macroprudential policy design relies on simplifying assumptions directed to make the problem analytically tractable. But, as in prior parts of this paper, the empirical formulations used in practice might be more complicated (and in some sense even more general) thus requiring adaptations that, in some cases, will also affect some of the properties uncovered below. In the spirit of keeping things simple and close to some of the previous extensions, assume that macroprudential policy involves M different dimensions, $j = 1, 2, ..M$ and that each dimension can now be associated with an intermediate objective I_j and a targeted policy tool z_j.

Assume further that intermediate objectives can be represented as linear functions of their targeted tools

$$I_j = \lambda_{0j} + \lambda_{1j} z_j, \tag{28}$$

where λ_{0j} is an autonomous component of intermediate objective and $\lambda_{1j} > 0$ measures the marginal impact of the targeted policy variable on the intermediate objective.[10] Then the baseline equations (2) and (3) could be reformulated as follows:

$$y_c = \alpha_c + \Gamma_c(I_1, I_2, ...I_M), \tag{29}$$

and

$$\bar{y} = \alpha + \Sigma_{j=1}^{M}\gamma_j z_j, \tag{30}$$

where Γ_c is an increasing and strictly concave function of the vector of intermediate objectives and $\gamma_j < 0$ so that, as in the baseline setup, macroprudential policies involve a trade-off: increasing policy z_j improves intermediate objective I_j but a the cost of reducing mean growth at the margin.[11]

As in the non-linear world described in the previous extension (on optimal policy mixes), all targeted policy variables activated at a strictly positive level at the optimum would satisfy a modified version of (7) and, consequently, (16), implying

$$\bar{y} - y_c = -\frac{1}{w}\frac{\frac{\partial \bar{y}}{\partial z_j}}{\frac{\partial \bar{y}}{\partial z_j} - \frac{\partial y_c}{\partial z_j}} = \frac{1}{w}\frac{1}{1 + \frac{\partial \Gamma_c}{\partial I_j}\lambda_j/(-\gamma_j)}. \tag{31}$$

Thus, as before, the optimal vector of targeted policies would feature equalization of the marginal growth-gap rates of transformation and, hence, *equalization of the marginal cost-effectiveness ratios* across all activated policy tools. Besides, the optimal gap between expected GDP growth and GaR would be decreasing in both the aversion to financial instability and such common ratio. The cost-effectiveness ratio of targeted policy j is the ratio between the marginal effectiveness of of the policy, that is, its marginal capability to improve GaR by affecting intermediate objective j ($\frac{\partial \Gamma_c}{\partial I_j}\lambda_j$) and the marginal cost of such policy in terms of mean growth ($-\gamma_j$).

Depending on the degree to which intermediate objectives may feature complementarity in their compound impact on GaR, as captured by the cross-derivatives of the function Γ_c, the setup with multiple intermediate objectives might imply increasing or decreasing the target gap as well as varying the optimal policy mix in response to changes in, for

---

[10]Without loss of generality, we impose a sign convention for I_j and z_j such that increasing I_j is *good* for financial stability (that is, increases y_c) and increasing the policy variable z_j is *good* for intermediate objective j (that is, increases I_j).

[11]Notice that, for simplicity and in contrast to the baseline specification, (29) and (30) do not explicitly contain any risk variable x. However, it would be trivial to introduce one or a vector of them affecting y_c and $\bar{y}$ linearly as in the baseline model. Additionally, one could also consider risk variables that affect the autonomous component λ_{0j} of each intermediate objective.

example, the autonomous component of one of the intermediate objectives. For example, in the simple case in which Γ_c were additively separable across intermediate objectives (so that they do not directly interact in affecting GaR), the policy response to a deterioration in the autonomous component of one objective j would be the tightening of policy across all intermediate objectives (so that $\frac{\partial \Gamma_c}{\partial I_{j'}}$ declines across all policy dimensions j' and the equality in (31) is restored at some higher target gap).

5.6 A discrete policy variable

Intuitively, if the policy variable is discrete and yet enters the problem as assumed in (2) and (3), the left hand side of the FOC in (7) must be replaced by its finite differences counterpart and its sign checked to discover if there are gains from increasing (or keeping increasing) the variable or, conversely, there might be gains from reducing it.

More formally, consider first the general case in which the policy variable can take N different values: $z \in \{z_1, z_2, ...z_N\}$ with $N \geq 2$. Let

$$\Delta W(x, z_i) = W(x, z_{i+1}) - W(x, z_i) \tag{32}$$

represent the welfare gain from increasing the discrete policy variable by one notch when starting from z_i. Using the definition of W in (5) and the expressions for (2) and (3), one can obtain

$$\Delta W(x, z_i) = \gamma(z_{i+1} - z_i) - \frac{w}{2}(\gamma_c - \gamma)^2(z_{i+1}^2 - z_i^2) + w(\gamma_c - \gamma)A(x)(z_{i+1} - z_i), \tag{33}$$

where $A(x) = (\alpha - \alpha_c) + (\beta - \beta_c)x > 0$. The first two terms in this expression are negative, reflecting the direct expected GDP cost of tightening macroprudential policy and the impact of such cost in reducing the gap between expected growth and GaR, which diminishes the marginal gains from further tightening. The third term is positive and increasing in the risk variable x and captures the gap reducing gains from tightening the policy. In a typical case, $\Delta W(x, z_i)$ will be positive at low values of i and turn negative at higher values of i, identifying the optimal policy as the highest i for which $\Delta W(x, z_i)$ is positive. Intuitively, as $A(x)$ is increasing, the optimal level of activation of the discrete policy will generally be higher for higher values of the risk variable x.

A particular case of interest in some applications is that in which the possible values of the policy variable are equally spaced (e.g. as when using a cumulative index of macroprudential policy actions). If one normalizes the scale of the variable to make the space between any two consecutive values to be one and sets $z_1 = 0$, then $z_i = i-1$ and one can use $z_{i+1}^2 - z_i^2 = 2i-1$

19

to write

$$\Delta W(x, z_i) = \gamma - \frac{w}{2}(\gamma_c - \gamma)^2(2i - 1) + w(\gamma_c - \gamma)A(x), \tag{34}$$

whose negative second term depends linearly on i reflecting, ceteris paribus, diminishing marginal welfare gains from activation of the discrete policy at higher and higher levels.

In the even more special case where the policy variable z is binary and can only take values 0 (inactive) or 1 (active), the welfare gain from activating the policy can be found setting $i = 0$ in (34):

$$\Delta W(x, 0) = \gamma - \frac{w}{2}(\gamma_c - \gamma)^2 + w(\gamma_c - \gamma)A(x), \tag{35}$$

whose interpretation is the same provided in the more general case.

In terms of Figure 1, the discreteness of the policy variable does not alter the indifference curves and the location of the "hypothetical" policy frontier that would emerge if z were continuous. The difference is that the effective frontier now only includes as many points on such hypothetical frontier as possible values z_i. Heuristically, it is still correct to think about the optimal policy as the one bringing the gap between expected growth and GaR as close as possible to the gap in (16) that would be targeted if z were a continuous variable.

6 Further discussion

6.1 What if the policy variable involves no trade-off?

Suppose the quantile regression methodology yields an estimate of γ equal to zero. In this case, under the remaining assumptions of the baseline model, z should be increased up to the point in which either the gap between expected growth and GaR is zero or the policy variable reaches its upper bound, whatever happens first. The first implication (being able to use the policy up to making GaR equal to expected growth) does not seem plausible or economically meaningful: it is too good to be true. In this case the emergence of $\gamma = 0$ in the estimation of the parameters of (3) may point to the existence of some relevant non-linearity (e.g., a negative effect which is observable only once z is large enough) that the linear specification fails to capture. In the field of macroprudential policy this can easily happen as many policies have not been historically used at all relevant ranges of activation, so identifying those negative effects in the data may be simply impossible.

Practical solutions to the problem may involve running non-linear specifications of (3) or, if the available data does not allows to capture the conjectured non-linearity, introducing the suspected missing cost of the policy using some auxiliary calculation. For instance, if

the policy variable is a borrower-based measure that has never being tried at a very high level but there are reasonable theoretical arguments to believe that its activation at too high levels may cause social unrest with, eventually, negative implications for welfare, one could introduce in the equation for expected GDP growth a negative term, increasing in z, that captures the estimated certainty-equivalent cost of the policy (expressed as a fraction of initial GDP). After such an adjustment, the design and assessment of macroprudential policy could proceed as indicated in prior sections

If the policy variable has a natural upper limit (e.g., is a binary or discrete variable measuring the quality of institutions such as, say, resolution regimes or policy coordination), then the implication that it should be activated at its maximum level may be meaningful and require no further adjustment in the analysis.

6.2 Country heterogeneity

In a multi-country environment, the empirical framework considered in this paper may involve country-specific versions of equations (2) and (3) as well as cross-country differences in the risk preference parameter w. Obtaining the former does not necessarily mean running quantile regressions country by country (which may suffer from severe data limitations) but, for instance, running panel quantile regressions allowing from country fixed effects or coefficients for the risk indicators or the policy variables that vary with some country-specific characteristics (e.g., variables intended to capture differences in the structure of countries' financial or legal systems). In the context of the "single country" baseline specification explored in this paper, these country differences can be thought as just having different values of the involved parameters and their implications can be easily extracted from the expressions for the policy rule and the target gap provided for the baseline case. In particular:

1. If countries structurally differ in aspects that only alter the intercepts α_c and α and/or the risk sensitivity parameters β_c and β in (2) and (3), then the target gap in (16) will not differ across countries. Yet, as reflected in the expressions for ϕ_0 in (10) and ϕ_1 in (11), their optimal policy rules may differ in intensity and risk responsiveness so as to accommodate their structural differences in each of these sets of parameters. For instance, a country with a larger value of $\alpha - \alpha_c$ (a larger "structural gap") will, other things equal, have to activate its macroprudential policy at a higher level (higher ϕ_0), while a country with a larger value of $\beta - \beta_c$ (a larger "gap vulnerability to risk") will have to be systematically more responsive to changes in the risk indicator x (higher ϕ_1).

2. If countries structurally differ in the effect of policy on GaR γ_c and/or expected growth γ, their target gap as well as the parameters of their optimal policy rule will differ. Specifically, countries featuring a hardest trade-off, as measured by the size of the marginal growth-gap rate of transformation (that is, the steepness of the policy frontiers depicted in Figure 1) will target a larger gap between expected growth and GaR and adapt their policy rules accordingly.

3. If countries structurally differ in their risk preferences as captured by w (a not very plausible source of heterogeneity under the microfoundation provided in the Appendix of this paper), then their target gap as well as the intercept ϕ_0 of their optimal policy rule will also differ. However, as previously mentioned when commenting on the determinants of ϕ_1, differences in w would not translate into a different responsiveness of their policies to changes in the risk variable x.

6.3 Interaction with other policies

The discussion so far has not explicitly dealt with the case in which policies other than macroprudential ones have an impact on expected growth and GaR. One immediate way to integrate them into the framework considered in this paper would be to add variables representing those policies in a vector version of the risk variable x. Under this reformulation, x would then account not only for risk variables in a narrow sense but also for other relevant elements of the economic situation at the time of designing macroprudential policy and that the macroprudential policymaker takes as given. Under this formulation (which resembles other treatments in which policies under the control of different authorities interact as in a non-cooperative game), the state at the time of designing macroprudential policy of other relevant non-macroprudential policies (e.g., monetary policy) would appear as part of x in vector versions of (2) and (3), and as a result in the macroprudential policy rule (9). The policy rule could then be interpreted as the macroprudential policy reaction function (reflecting the reaction of macroprudential policy to the current settings of other policies).

A more general discussion covering the issue of optimal policy coordination would require further extensions. For instance, to consider optimal coordination with monetary policy, the objective function W might have to include terms reflecting goals of such policy, such as price stability, that might not be fully reflected in the terms currently included in (5).[12] A policymaker optimizing on the two policies at the same time would in that case treat the

---

[12]See Cecchetti and Kohler (2014) for an example that considers coordination between conventional monetary policy and capital regulation in a related reduced-form setup.

non-macroprudential policy under consideration as an element of a vector version of the policy variable z giving rise to issues similar to those discussed in Sections 5.3 and 5.4 when considering multiple macroprudential policy tools.

6.4   Reformulation in terms of growth-given-stress

Interestingly, in a normal distribution the distance between the mean and the c-quantile is proportional to the distance between the mean and the expectation of the random variable conditional on being below the c-quantile. Section A.3 of the Appendix uses this property to show that, under the baseline assumption that the GDP growth rate is normally distributed, the welfare criterion in (5) can be re-expressed in terms of growth-given-stress (that is, the expectation of the growth rate conditional on being below the c-quantile of the growth rate), retaining the microfoundation provided in Section A.1 of such Appendix. Consequently, the constant target gap property of the optimal policy rule in (16) could also be expressed in terms of the distance between expected growth and growth-under-stress. So, under normality, formulating the macroprudential policy problem and the assessment of macroprudential stance on the basis of GaR or growth-given-stress would make no difference.

7   Concluding remarks

Using the concept of GaR in the measurement of the downside risks that macroprudential policy aims to address opens very interesting avenues for the use of empirical quantitative models for the design of macroprudential policies and for the development of concrete notions of macroprudential policy stance. The setup allows to explicitly consider, under a similar modeling methodology, the effects of risk and policy variables on expected GDP growth (arguably, a succinct measure of what other macroeconomic policies care about) and the risk of sufficiently adverse GDP growth outcomes (arguably, a promising concrete measure of what macroprudential policy cares about). This paper has explored the foundations for the design and assessment of macroprudential policies using this setup.

The paper has started with a very stylized description of the setup in the context of its implementation using the outcome of a quantile regression approach. A welfare criterion for the design of the optimal policies has been proposed that can be microfounded as consistent with the maximization of the expected utility of a representative agent in some contexts. The properties of the optimal policies have been explored in the basic setup as well as in several extensions and modifications covering cases with non-linearities in the impacts of policy variables and risk variables on the relevant outcomes, multiple policy variables, and discrete

policy variables. An important extension has shown the compatibility of this framework with the view that macroprudential policy involves various well-identified intermediate objectives each of which can be associated with one or a subset of targeted policy tools. Additional discussions have dealt with policies that seem to involve no trade-off between mean growth and GaR, the treatment of country heterogeneity, and the possibility of reformulating the analysis around the concept of growth-given-stress rather than GaR.

Under the postulated representation of preferences, the policy design problem yields a quantitative-based policy target and a metric for the assessment of policy stance similar to that of other macroeconomic policies. The main challenges ahead for the applicability of this framework are more empirical and political than conceptual. On the empirical side. the main challenge resides at the consistent and precise enough estimation of the causal effects of risk and policy variables on the relevant moments (mean and GaR) of the growth distribution. Properly detecting relevant non-linearities and interactions between policies is also important. Absent proper estimates of the relevant parameters and relationships a mechanical application of this framework could produce misguided policy advice. So the framework will develop at the speed with which data on the applied policies accumulates and econometric efforts succeed in providing reliable estimates of their effects on growth outcomes.

On the political side, once data and estimation provide a reliable description of the policy trade-offs, the main challenge is at defining the society's aversion for financial instability on which optimal policies should be based. Given the uncertainty on the relevant parameters implied by the empirical challenges, policymakers many need to be guided on how to expand the type of framework sketched in this paper to account for model uncertainty (that is, for the imperfect knowledge of the specification and parameters of the relevant quantile regressions) and the potential policy mistakes that could stem from such uncertainty.

References

Adrian, T., N. Boyarchenko, and D. Giannone (2019) "Vulnerable growth," *American Economic Review*, 109, 1263-89.

Adrian, T., F. Grinberg, N. Liang, and S. Malik (2018) "The term structure of Growth-at-Risk," International Monetary Fund Working Paper 18/180.

Aikman, D., J. Bridges, S. Hacioglu, C. O'Neill, and A. Raja (2019) "Credit, capital and crises: a GDP-at-Risk approach," Bank of England Staff Working Paper 724.

Caldera-Sánchez, A., and O. Röhn (2016) "How do policies influence GDP tail risks?," OECD Economics Department Working Papers 1339.

Cecchetti, S. (2008) "Measuring the macroeconomic risks posed by asset price booms," in J. Campbell, editor, *Asset Prices and Monetary Policy*, University of Chicago Press for NBER.

Cecchetti, S., and M. Kohler (2014) "When capital adequacy and interest rate policy are substitutes (and when they are not)," *International Journal of Central Banking*, 10, 205-232.

Cecchetti, S., and H. Li (2008) "Measuring the impact of asset price booms using quantile vector autoregressions," Brandeis University, mimeo.

Cecchetti, S., and J. Suarez (2021) "On the stance of macroprudential policy," Brandeis University, mimeo.

Chavleishvili, S., R. Engle, S. Fahr, M. Kremer, S. Manganelli, and B. Schwaab (2020) "The risk management approach to macro-prudential policy," European Central Bank, mimeo.

De Nicolo, G., and M. Lucchetta (2017) "Forecasting tail risks," *Journal of Applied Econometrics*, 32, 159-170.

Duprey, T., and A. Ueberfeldt (2018) "How to manage macroeconomic and financial stability risks: A new framework," Bank of Canada Staff Analytical Note 11.

Fishburn, P. (1977) "Mean-risk analysis with risk associated with below-target returns," *American Economic Review*, 67, 116-126.

Galán, J. (2021) "The benefits are at the tail: uncovering the impact of macroprudential policy on GDP growth," *Journal of Financial Stability*, forthcoming.

Kahneman, D., and A. Tversky (1979) "Prospect theory: An analysis of decision under risk," *Econometrica*, 47, 263–291.

Kilian, L., and S. Manganelli (2008) "The central banker as a risk manager: Estimating the Federal Reserve's preferences under Greenspan," *Journal of Money, Credit and Banking*, 40, 1103-1129.

Prasad, A., S. Elekdag, P. Jeasakul, R. Lafarguette, A. Alter, A. X. Feng, and C. Wang (2019) "Growth at Risk: Concept and application in IMF country surveillance," International Monetary Fund Working Paper 19/36.

Svensson, L. (2003) "Optimal policy with low-probability extreme events," in *Macroeconomics, Monetary Policy, and Financial Stability – A Festschrift for Charles Freedman*, Proceedings of a conference held by the Bank of Canada, Ottawa, June 2003, 79-104.

Appendix
Microfoundations of the GaR-based welfare criterion

A.1 CARA preferences and normally distributed growth rates

Let Y denote GDP and let y describe the implied (geometric) GDP growth rate relative to a benchmark level Y_0 so that

$$Y = (1 + y)Y_0. \tag{36}$$

Suppose also that there is a representative agent whose preferences for GDP levels are represented by a utility function $U(Y)$ with a local coefficient of *absolute* risk aversion $\lambda(Y_0)$ at $Y = Y_0$ and that the utility function can be (locally) described as one exhibiting CARA with parameter $\lambda(Y_0)$, so that

$$U(Y) = -\exp(-\lambda(Y_0)Y). \tag{37}$$

Using (36), we can write

$$U(Y) = -\exp(-\lambda(Y_0)Y_0(1 + y)) = -\exp(-\lambda(Y_0)Y_0)\exp(-\lambda(Y_0)Y_0 y). \tag{38}$$

For fixed Y_0, since affine monotonic transformations of a utility function will represent exactly the same preferences, we can replace $U(Y)$ with

$$u(y) = -\exp(-\lambda(Y_0)Y_0 y) = -\exp(-\rho_0 y), \tag{39}$$

where $\rho_0 = \lambda(Y_0)Y_0$ describes the agent's coefficient of *relative* risk aversion at Y_0. Thus, this utility function describes CARA preferences directly on the growth rate y but the parameter ρ_0 in such specification measures the relative risk aversion of the agent (in terms of her preferences for GDP levels) at the initial GDP level Y_0.

Suppose now that GDP growth is normally distributed, so $y \sim N(\bar{y}; \sigma_y^2)$. From the well-known properties of a normal distribution, the moment generating function of the distribution of y is then

$$M(t) = E(\exp(ty)) = \exp(\bar{y}t + \frac{1}{2}\sigma_y^2 t^2) \tag{40}$$

for any t. So, in particular,

$$M(-\rho_0) = E(\exp(-\rho_0 y)) = \exp(-\rho_0\bar{y} + \frac{1}{2}\rho_0^2\sigma_y^2). \tag{41}$$

Hence, from (39) and (41), we can write the agent's expected utility as

$$E[u(y)] = -E[\exp(-\rho_0 y)] = -\exp(-\rho_0\bar{y} + \frac{1}{2}\rho_0^2\sigma_y^2). \tag{42}$$

And, since monotonic transformations of expected utility will represent exactly the same preferences, such preferences can be equivalently described by the (indirect) utility function

$$v = \bar{y} - \frac{\rho_0}{2}\sigma_y^2. \tag{43}$$

that is, a simple linear expression in the mean $\bar{y}$ and the variance σ_y^2 of the growth rate y.

The Growth-at-Risk (GaR) for a given confidence level c is the c-quantile of the probability distribution of y, that is, the value y_c such that

$$\Pr(y \le y_c) = c. \tag{44}$$

By the properties of normal distributions, $(y - \bar{y})/\sigma_y$ is a standard normal random variable, $N(0,1)$. Letting $\Phi(\cdot)$ be the cumulative distribution function of a standard normal, we can write

$$\Pr(y \le y_c) = c \Leftrightarrow \Pr((y - \bar{y})/\sigma_y \le (y_c - \bar{y})/\sigma_y) = c \Leftrightarrow \Phi((y_c - \bar{y})/\sigma_y) = c. \tag{45}$$

Solving for y_c in the last expression yields

$$y_c = \bar{y} + \sigma_y \Phi^{-1}(c). \tag{46}$$

Alternatively, solving for σ_y yields

$$\sigma_y = \frac{y_c - \bar{y}}{\Phi^{-1}(c)}, \tag{47}$$

which plugged into (43) leads to the indirect utility function

$$v(\bar{y}, y_c; \rho_0, c) = \bar{y} - \frac{\rho_0}{2(\Phi^{-1}(c))^2}(\bar{y} - y_c)^2, \tag{48}$$

which expresses the agent's expected utility as a function of expected growth, GaR at a confidence level c, the relative risk aversion coefficient of the agent at the initial level of GDP ρ_0, and the confidence level c.

Hence, maximizing a welfare criterion of the form

$$W = \bar{y} - \frac{w}{2}(\bar{y} - y_c)^2, \tag{49}$$

as assumed in the main text, would be equivalent to the maximization of the expected utility of the representative agent for

$$w = \frac{\rho_0}{(\Phi^{-1}(c))^2}. \tag{50}$$

For instance, for $c = 0.05$, one has $\Phi^{-1}(c) = -1.6449$, so with a coefficient $\rho_0 = 2$ of relative risk aversion at Y_0, both criteria would coincide under $w = 2(1.6449)^{-2} = 0.7392$.

Quite intuitively, the policy maker's preference for financial stability should increase with the agent's relative risk aversion parameter ρ_0 as well as, for any $c < 0.5$, with the level of confidence c at which GaR is calculated.[13]

A.2 Modeling GaR vs. growth volatility and departing from normality

Under the normality assumption sustaining the interpretation of the welfare criterion W as consistent with expected utility maximization, modeling a lower quantile such as y_c and expected growth $\bar{y}$ is not different from modeling the standard deviation and the mean of the growth rate and focusing on a welfare criterion that directly depends on those moments of the growth distribution.

Moreover, under normality, if expected growth is determined as in (3) and the standard deviation of the growth rate is linear in x and z, say

$$\sigma_y = \alpha_\sigma + \beta_\sigma x + \gamma_\sigma z, \tag{51}$$

then (46) implies that (51) is exactly compatible with the specification of y_c in (2) if and only if $\alpha_c = \alpha + \Phi^{-1}(c)\alpha_\sigma$, $\beta_c = \beta + \Phi^{-1}(c)\beta_\sigma$, and $\gamma_c = \gamma + \Phi^{-1}(c)\gamma_\sigma$, where for $c < 0.5$ we have $\Phi^{-1}(c) < 0$. So, for instance, the prior assumption that the policy variable has a positive effect on y_c ($\gamma_c > 0$) and a negative effect on $\bar{y}$ ($\gamma < 0$) would require that it also has a sufficiently large negative impact on σ_y ($\gamma_\sigma < -\gamma/\Phi^{-1}(c) < 0$).

While having the capacity to structurally interpret the analysis in the main text under the assumption of normality as exactly compatible with expected utility maximization is reassuring, the normal case would not justify a strict preference for the quantile regression approach. A quantile regression approach to the analysis of macroprudential policies is typically defended on the grounds that there are variables whose impact on extreme low quantiles of the growth distribution is empirically detectable, while its impact on the standard deviation of the growth rate (or on high quantiles of the growth distribution) might not be (or at least not so clearly). For instance, it is likely that empirical measures of GDP volatility are dominated by what happens at business cycle frequencies, while what happens at a sufficiently low growth quantile may better capture the impact of infrequent financial crises.

However, representing the world in which lower quantiles are disproportionately affected by one variable or infrequent discrete events have non-linear implications for growth implies departing from the normality assumption and, hence, from the setup in which the interpretation of the welfare criterion in expected utility terms is exactly valid. In other words, while the normal world provides a useful benchmark to help connect the preference for financial stability reflected into the welfare criterion W with a standard way of representing agent's

---

[13]Notice that, for $c < 0.5$, $\Phi^{-1}(c) < 0$ and approaches zero as c increases, so $(\Phi^{-1}(c))^2$ is decreasing in c.

preferences in economics, it is probably not the most practically relevant one. Box D in Cecchetti and Suarez (2021) describes a number of simulation exercises in which GDP growth is drawn from empirically-motivated non-normal distributions and examines the accuracy with which a GaR-based criterion such as W approximates the true expected utility (measured in certainty equivalent terms). The message from those simulations is that the GaR-based metrics provides a reasonably good approximation to the expected-utility-based metrics even when the growth distribution deviates substantially from normality, as well as under CRRA preferences.

Additionally, as discussed in the main text, in the non-normal world, one might interpret W as a heuristic representation of the preferences of a policy maker who cares about the gap $\bar{y} - y_c$, for a suitably low value of c, rather than, say, the standard deviation of GDP growth, because of some form of loss aversion. Under this perspective, the focus on the trade-off between maximizing $\bar{y}$ and minimizing the gap $\bar{y} - y_c$ could reflect that the policy maker cares more about the relative output losses incurred at the low tail of the growth rate distribution than the potentially offsetting (in expected terms) relative output gains obtained at the high tail.

A.3 Reformulation using a growth-given-stress criterion[14]

Define the Growth-given-Stress (GgS) for a given reference probability c as the expected value of the GDP growth rate y conditional on such rate being lower than the c-quantile of its distribution y_c, that is

$$GgS_c = E(y \mid y \leq y_c). \tag{52}$$

When y is a normal random variable, GgS_c is just the mean of an upper-truncated normal random variable with truncation point at y_c. The well-known expression for such mean implies

$$GgS_c = \bar{y} - \frac{\phi\left(\frac{y_c - \bar{y}}{\sigma_y}\right)}{\Phi\left(\frac{y_c - \bar{y}}{\sigma_y}\right)}\sigma_y,$$

where $\phi(\cdot)$ is the density function of standard normal. But, since y_c is the c-quantile of the distribution of y, the term $(y_c - \bar{y})/\sigma_y$ can be written as $\Phi^{-1}(c)$. This allows us to write

$$GgS_c = \bar{y} - \frac{\phi\left(\Phi^{-1}(c)\right)}{c}\sigma_y.$$

Now, using (47) to substitute for σ_y and re-arranging, one can express

$$\bar{y} - GgS_c = \frac{-\phi\left(\Phi^{-1}(c)\right)}{c\Phi^{-1}(c)}\left(\bar{y} - y_c\right),$$

---

[14]I thank Steve Cecchetti for making me notice the possibility of this reformulation.

where, for given c, the ratio $-\phi\left(\Phi^{-1}(c)\right)/\left(c\Phi^{-1}(c)\right)$ is a proportionality constant (which is positive for $c < 0.5$).

In words, when the growth rate y is normally distributed the gap between expected growth and GgS is proportional to the gap between expected growth and GaR. Therefore, maximizing the welfare criterion W specified in (5) would be equivalent to maximizing a similar linear-quadratic criterion whose quadratic term contains the square of the distance between expected growth and GgS and where the instability aversion parameter w is replaced by

$$w_{GgS} = \left(\frac{c\Phi^{-1}(c)}{\phi\left(\Phi^{-1}(c)\right)}\right)^2 w.$$

Such criterion would thus have the same microfoundation as the one provided in Section A1 of this Appendix for W. Under such microfoundation, the parameter w_{GgS} would become, using (50),

$$w_{GgS} = \frac{c^2\rho_0}{\phi(\Phi^{-1}(c))}$$

The optimal policy rule resulting from solving the baseline policy problem under the GaR-based welfare criterion would be equivalent to the one emerging under the equivalent GgS-based criterion, and would also satisfy the constant target gap property in (16). Such property could be translated into targeting a gap between expected growth and GgS given by

$$\bar{y} - GgS_c = \frac{1}{w_{GgS}}\frac{1}{1+\gamma_c/(-\gamma)}.$$