Discussion of "Macroprudential Policy under Uncertainty"*

François Gourio
Federal Reserve Bank of Chicago

1. Introduction

Following the global financial crisis of 2007–08, policymakers around the world have made it a priority to avoid another occurrence of financial crisis. New, tighter regulations of banks and other financial intermediaries, but also mortgage and other credit markets, hold the promise of reducing both the likelihood and severity of a financial crisis. More controversially, some countries have adopted time-varying macroprudential tools, such that these regulations can be changed according to economic conditions.[1]

Of course, setting the appropriate level of macroprudential "tightness" is challenging, even more so if it has to be varied over time. Policymakers face all the usual difficulties in measuring in real time the relevant economic and financial variables, weighting complicated trade-offs to arrive at an optimal policy, then communicating it and implementing it. On top of that, the effect of the tools is extremely uncertain because these tools are new and untested. To be sure, there is accumulating evidence on the relation between credit, growth, and financial risk (e.g., Schularick and Taylor 2012). However, this evidence is still preliminary and the estimated

*The views expressed here are mine only and do not necessarily reflect those of the Federal Reserve Bank of Chicago or the Federal Reserve System. Author contact: Federal Reserve Bank of Chicago, 230 S. Lasalle St, Chicago IL 60604. E-mail: francois.gourio@chi.frb.org.

[1]Perhaps even more controversially, some policymakers consider using monetary policy tools such as the level of short-term nominal interest rates to affect credit conditions and hence ultimately the risk of financial crisis. See Stein (2013), Ajello et al. (2016), Svennson (2016), and Gourio, Kashyap, and Sim (2017) for discussion and analysis.

relationships may be unstable.[2] The shift in policy regime also means that existing empirical evidence may be less relevant because of the Lucas critique.

In this context, the current paper is a welcome and timely addition. The key contribution is to explain very clearly under which circumstances policymakers should take into account these various forms of uncertainty when formulating policy. The paper builds on a well-established normative literature that studies how uncertainty should affect monetary policy. The starting point for this literature is the certainty equivalence principle, which states that if an agent (the policymaker) is choosing an action to minimize a quadratic expected loss subject to linear constraints and additive shocks, then higher uncertainty about the effect of shocks has no effect on the optimal action: the policymaker takes the same action as if the shock was known to be zero. Of course, higher uncertainty is bad: it increases the expected loss, but it does not affect the optimal action.

Following up on the previous literature, the paper shows that any deviation from these assumptions can generate a deviation from certainty equivalence, i.e., optimal policy is affected by the amount of uncertainty. Clearly, if the loss function is not quadratic—for instance, it is asymmetric—or if the economy is not described by a set of linear equations, the certainty equivalence principle will be violated, though the direction of the bias will depend of course on the exact asymmetry. Also, following Brainard (1967), uncertainty about parameters can lead to attenuation bias; however, as has been recognized for some time, this result is not general and it is easy to build counterexamples where parameter uncertainty lead to a more aggressive policy (see, for instance, Söderström 2002).

Many results discussed in the paper are drawn from literature, but the authors add a very clear exposition and a thoughtful discussion of the practical relevance for macroprudential policy. The authors also propose a new "signaling channel" through which the regulator may use his macroprudential tool not only because it

[2]There is a long history of breakdown of empirical relationships between macroeconomic variables and financial variables, and they may not be informative in real time (e.g., Edge and Meisenzahl 2011), due in part to financial innovations that make consistent financial measurement challenging.

affects the actual risk but also because agents infer from its use that the central bank sees a high level of financial fragility. This allows policy to be more powerful and hence be used more sparingly. An interesting question is whether this allows the policy to be also less distortive (i.e., are signals "cheaper" than actions?).

One potential limitation of the paper is that it studies the policy problem in a simple static linear-quadratic framework. The advantage is the clarity, but it comes at the cost of abstracting from several issues. First, it is not quite clear if some results apply to the *level* of policy variable or to the *change*. For instance, the policy question may be whether to lower the capital requirement from 12 percent to 10 percent. In this context, what does "inaction" or "attenuation bias" mean? Getting a smaller ratio may be less active, but the change in itself is an action.

Second, the static framework abstracts from issues of expectations that may be important, such as these found in the study of monetary policy. For instance, suppose that a regulator adopts a more reactive macroprudential policy rule. This would create expectations by banks of changes in capital requirements in the future. Since banks do not like to have to raise capital quickly, this rule would likely change the behavior of banks which might keep larger buffers (above the capital requirement). These kind of implications are absent from the analysis.

My third issue with the framework has to do with the linear-quadratic assumptions. This is a natural first step, and these assumptions have been commonly used to study monetary policy. However, the framework may be substantially more appropriate for the stabilization problem inherent in monetary policy—i.e., trading off the variance of output against the variance of inflation. It might not adequately capture the trade-offs relevant for prudential policy. For instance, very high financial stability is assumed to be "bad," which on its face seems difficult to understand. I believe it is more natural to think of macroprudential policy as involving a trade-off between the likelihood of crisis with a steady-state cost to the level of output in the economy. Hence, instead of a variance-variance trade-off, one has to consider a mean-variance trade-off. In the rest of this discussion I provide a few simple examples of how one might start thinking about this trade-off.

2. Macroprudential Policy as a Mean-Variance Trade-Off

I first discuss the case with known parameters, then consider the case with unknown parameters.

2.1 The Case with Known Parameters

For simplicity, start with a static framework. Agents have constant relative risk aversion (CRRA) utility over aggregate consumption,

$$U(C) = \frac{C^{1-\gamma}}{1-\gamma}$$

and aggregate consumption is uncertain. To start, assume that it is log-normally distributed:

$$\log C \hookrightarrow N\left(\mu - \frac{\sigma^2}{2}, \sigma^2\right)$$

where μ is the log of the mean of consumption, i.e., $E(C) = e^{\mu}$.[3] Suppose the regulator (government) can take some action that will affect both the mean and the variance of consumption, such as tighter bank capital requirements. Such a measure may reduce variance at the risk of lowering the mean. One can capture this in a reduced-form way by assuming a relationship

$$\mu = \mu_0 + \mu_1 \sigma,$$

where $\mu_1 < 0$ captures the slope of the trade-off.

If the parameters (μ_0, μ_1) describing the trade-off are known, the optimal policy involves choosing a risk that maximizes expected utility, i.e.,

$$\max_{\sigma} E\left(U(C)\right).$$

[3] I incorporate the standard Jensen adjustment, i.e., the mean of log consumption is $\mu - \sigma^2/2$, which ensures that the level of risk σ does not affect the mean of consumption.

It is easy to calculate this expected utility, which turns out to be equal to

$$\max_{\sigma} U\left(\mu_0 - \mu_1\sigma - \gamma\frac{\sigma^2}{2}\right),$$

so that the term inside is the certainty equivalent of consumption. Maximizing expected utility involves then maximizing the certainty equivalent:

$$\max_{\sigma}\left(\mu_0 - \mu_1\sigma - \gamma\frac{\sigma^2}{2}\right).$$

The solution is characterized by the first-order condition which leads to

$$\sigma^* = \frac{-\mu_1}{\gamma},$$

so that higher risk aversion γ will lead to a choice of a lower variance σ^*, while a steeper trade-off (a higher absolute value of μ_1) will lead to a higher risk since reducing risk is more costly.

2.2 The Case with Unknown Parameters

What happens if, on the other hand, the parameters (μ_0, μ_1) are not known with certainty? I will solve this in the simple case where the decisionmaker has a prior normal distribution over each of these parameters. It turns out that the Brainard logic presented by the authors applies in this context,[4] but as we will see, the conclusion might have a slightly different interpretation.

First, to make sure that an increase in uncertainty about μ_0 or μ_1 has no direct effect on expected consumption, I will assume that

$$\mu = \widetilde{\mu_0} + \widetilde{\mu_1}\sigma + k(\sigma),$$

[4]Essentially, the combination of CRRA utility, log-normally distributed endowment, and normal prior yields a quadratic problem similar to the one studied by the authors.

where $k(\sigma)$ is chosen so that $E(C) = e^{E(\mu_0)-E(\mu_1)\sigma}$ does not depend on $Var(\mu_0)$ or $Var(\mu_1)$.[5] Log consumption is still assumed to be distributed with mean $\mu - \frac{1}{2}\sigma^2$ and variance σ^2.

Calculating expected utility (where the expectation is taken now both over the shock and the prior distribution over the parameters) yields

$$E\left(U(C)\right) = U\left(E(\mu_0) - \frac{\gamma}{2}Var(\mu_0) - E(\mu_1)\sigma - \gamma\frac{\sigma^2}{2}Var(\mu_1) - \gamma\frac{\sigma^2}{2}\right)$$

so that the optimal policy problem is to maximize the certainty equivalent:

$$\max_{\sigma} E(\mu_0) - \frac{\gamma}{2}Var(\mu_0) - E(\mu_1)\sigma - \gamma\frac{\sigma^2}{2}Var(\mu_1) - \gamma\frac{\sigma^2}{2}$$

with first-order condition leading to the optimal choice:

$$\sigma^* = \frac{-E(\mu_1)}{\gamma(1 + Var(\mu_1))}.$$

We see that the same two conclusions as in the Brainard model hold here: first, uncertainty about μ_0 is irrelevant for the optimal decision—the term $Var(\mu_0)$ does not appear in the optimal policy choice σ^*, but it does reduce the expected utility. Second, uncertainty about μ_1 leads to a choice of a lower variance σ^* than the policymaker would choose if he had full information about the parameters. One might think that is a standard Brainard result, and indeed it is, except for the interpretation—a low σ^* means, in practice, a very tight macroprudential policy, so the policymaker is actually using his tools aggressively despite the high uncertainty with which they affect the average consumption of the economy. But as discussed above, this reflects that the Brainard result does not hold under all conditions—it depends if a more aggressive policy increases or reduces risk.

[5]This requires that

$$k = -\frac{1}{2}Var(\mu_0) - \frac{1}{2}\sigma^2 Var(\mu_1),$$

simply reflecting the standard Jensen adjustment.

The result above relies on many assumptions, including the functional form linking μ and σ, as well as the utility function and the log-normal assumption. It should be clear that altering these assumptions could easily lead to a reversal of the result.

2.3 Modeling Financial Crises as Rare Disasters

The modeling so far has assumed that macroprudential policy aims to stabilize the economy by reducing the variance of consumption. Macroprudential policy is, however, mostly focused on reducing the likelihood or severity of rare financial crises. One might worry then that the modeling adopted above, while tractable, does not capture the key trade-offs appropriately.

To illustrate this, I consider the kind of processes studied in Barro (2006) or Gourio (2012): consumption may have either a high level $\overline{C}$ or a low level $\overline{C}(1-b)$. The first case arises with probability $1-p$ and the second one with probability p. The scalar p can be thought of as the probability of a financial crisis, and b measures the share of consumption lost in a financial crisis. In this case, it is easy to see that expected utility is

$$EU(C) = \frac{\overline{C}^{1-\gamma}}{1-\gamma}\left(1 - p + p(1-b)^{1-\gamma}\right).$$

If one assumes that there is a trade-off between the level of consumption and the probability of crisis, for instance,

$$\log \overline{C} = d_0 + d_1 p,$$

where the parameters (d_0, d_1) reflect the trade-off between the level of output and the probability of financial crisis, then it is straightforward to obtain the optimal p. One could also do the same calculation if there is uncertainty about d_1. Clearly, the functional forms and assumptions will drive the result. The general spirit is quite similar to that of the log-normal example. But the quantitative results might be quite different, as the welfare cost of avoiding rare and large shocks can be quite larger than smoothing out small fluctuations (e.g., Barro 2009; Gourio, Kashyap, and Sim 2017).

2.4 Generalization to Infinite Horizon

For simplicity I have focused on static calculations in these examples, and I have used a simple power utility. The goal of this section is simply to clarify that this comes at no cost of generality, because it is straightforward to generalize these results to an infinitely lived agent with Epstein-Zin preferences, at least as long as consumption growth is iid. To see this, consider a representative agent in which utility (value) is defined recursively as

$$V_t = \left((1-\beta) C_t^{1-1/\psi} + \beta E_t \left(V_{t+1}^{1-\gamma} \right)^{\frac{1-1/\psi}{1-\gamma}} \right)^{\frac{1}{1-1/\psi}},$$

where ψ is the elasticity of intertemporal substitution to consumption and γ measures risk aversion. Assume that this consumer has an initial consumption C_0 and faces macroeconomic risk in the form of a unit-root iid consumption process:

$$\Delta \log C_{t+1} = X_{t+1},$$

where X_{t+1} is an iid random variable. For instance, X_{t+1} could be normally distributed with mean $\mu - \frac{\sigma^2}{2}$ and variance σ, the standard log-normal case. But one could also allow X_{t+1} to take the form of rare financial crises, so that $X_{t+1} = \mu$ with probability $1 - p$ and $X_{t+1} = \mu - b$ with probability p.

To find the optimal policy, we must first solve for welfare. It is easy to see in this model that the value-consumption ratio V_t/C_t is constant and given by the solution to

$$v^{1-\frac{1}{\psi}} = 1 - \beta + \beta v^{1-\frac{1}{\psi}} E_t \left(X_{t+1}^{1-\gamma} \right)^{\frac{1-1/\psi}{1-\gamma}}$$

and hence is determined by the moment $E_t(X_{t+1}^{1-\gamma})$, which is the one we studied in the static model. Hence our conclusions should apply in this framework as well.

3. Conclusion

The paper is a clean and lucid analysis of how uncertainty should affect macroprudential policy. It is hard to draw strong conclusions given the various channels and effects highlighted by the authors. Moving beyond the qualitative results will require a credible DSGE model that incorporates financial crises, makes their probability and severity affected by macroprudential policy, and also incorporates the costs of tighter regulation. Such a model would address the key trade-off created by macroprudential policy, namely that lower volatility (lower financial crisis probability) comes at a cost of less capital accumulation and less production.

To conclude, one should also keep in mind that the design of macroprudential policy cannot be isolated from that of the other policies (monetary and fiscal). A natural approach is to let macroprudential policy focus on financial stability and let the other policies offset the indirect effect of macroprudential policy on output and inflation. This might create an additional constraint on monetary and fiscal policies, and transfer the risk from the regulatory to the monetary authority.[6]

References

Ajello, A., T. Laubach, D. López-Salido, and T. Nakata. 2016. "Financial Stability and Optimal Interest-Rate Policy." Working Paper, Board of Governors of the Federal Reserve System.

Barro, R. 2006. "Rare Disasters and Asset Markets in the Twentieth Century." *Quarterly Journal of Economics* 121 (3): 823–66.

———. 2009. "Rare Disasters, Asset Prices, and Welfare Costs." *American Economic Review* 99 (1): 243–64.

Brainard, W. 1967. "Uncertainty and the Effectiveness of Policy." *American Economic Review* 57 (2): 411–25.

Edge, R., and R. Meisenzahl. 2011. "The Unreliability of Credit-to-GDP Ratio Gaps in Real Time: Implications for Countercyclical Capital Buffers." *International Journal of Central Banking* 7 (4, December): 261–98.

[6]For instance, in the spirit of Evans et al. (2015), a tight macroprudential policy might make it more likely that monetary policy is constrained (now or in the future) by the zero lower bound on nominal interest rates, increasing the risk of not reaching its targets and increasing uncertainty for the economy more generally.

Evans, C., J. Fisher, F. Gourio, and S. Krane. 2015. "Risk Management near the Zero Lower Bound." *Brookings Papers on Economic Activity* (Spring): 141–219.

Gourio, F. 2012. "Disaster Risk and Business Cycles." *American Economic Review* 102 (6): 2734–66.

Gourio, F., A. Kashyap, and J. Sim. 2017. "The Tradeoffs in Leaning Against the Wind." Working Paper, Federal Reserve Bank of Chicago.

Schularick, M., and A. Taylor. 2012. "Credit Booms Gone Bust: Monetary Policy, Leverage Cycles, and Financial Crises, 1870–2008." *American Economic Review* 102 (2): 1029–61.

Söderström, U. 2002. "Monetary Policy with Uncertain Parameters." *Scandinavian Journal of Economics* 104 (1): 125–45.

Stein, J. 2013. "Overheating in Credit Markets: Origins, Measurement, and Policy Responses." Speech given at the "Restoring Household Financial Stability after the Great Recession: Why Household Balance Sheets Matter" research symposium sponsored by the Federal Reserve Bank of St. Louis, held in St. Louis, Missouri, February 7. Available at https://www.federalreserve.gov/newsevents/speech/stein20130207a.htm.

Svensson, L. E. O. 2016. "Cost-Benefit Analysis of Leaning Against the Wind." NBER Working Paper No. 21902.