"Progressive Taxation and Monetary Policy in Australia"

Discussion remarks by Sushant Acharya (University of Melbourne)

1. Introduction

This paper is part of a larger effort by the Reserve Bank of Australia to develop a Heterogeneous Agent New Keynesian (HANK) model tailored to the specifics of the Australian economy. The plan is to eventually integrate this model into the suite of models used as inputs into policymaking at the Bank. The paper uses a variant of the aforementioned HANK model to study how the progressivity of the tax system affects the tradeoffs facing monetary policy. In particular, the paper focuses on three dimensions of how tax progressivity interacts with monetary policy. First, it examines how a greater progressivity of the tax system affects the natural rate of interest (r^*) through its effect on households' precautionary savings behavior. Second, it explores how progressive taxation changes the sensitivity of output to monetary policy shocks. Finally, the paper also explores how a monetary policy shock differentially impacts the welfare of households across the wealth and income distribution.

In my discussion, I focus on three issues. First, I assess whether the estimated quantitative effects are meaningful for policy. Second, I interpret the mechanisms that drive these results and discuss their robustness. Third, I reflect on the broader methodological challenges of using HANK models to address practical policy questions.

2. Main Takeaways

2.1 Effect on $r^\star$

The first exercise in the paper studies how an increase in tax progressivity can affect the natural rate of interest, which is measured in the model as the steady-state real interest rate. Now, why should tax progressivity matter for the natural rate of interest? The logic of HANK models implies that the natural rate of interest is determined by the households' demand for savings. Holding all else constant, any force which *reduces* the demand for savings will tend to *increase* the natural rate. An increase in

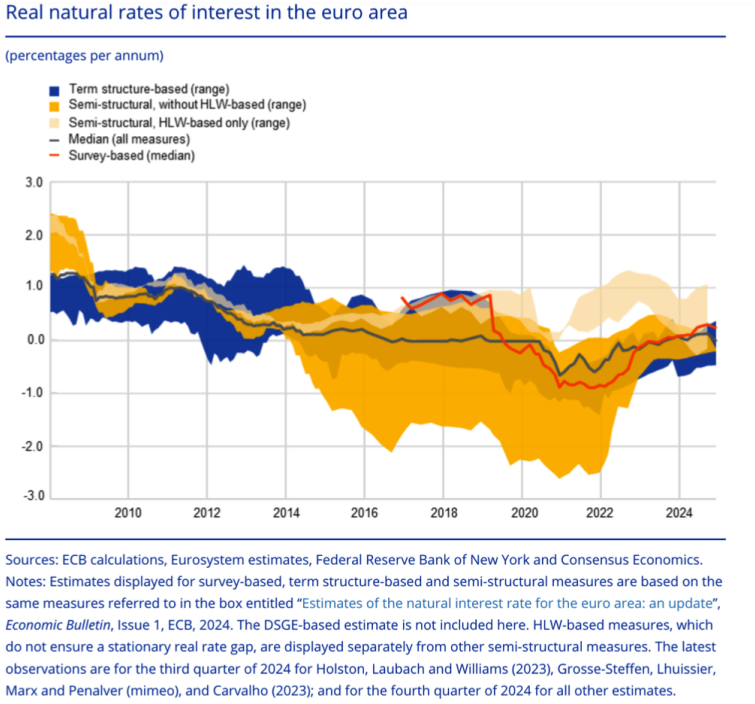

Figure 1: Uncertainty in $r^\star$ estimates in the Euro area

tax-progressivity does precisely this. A more progressive tax system provides insurance to households as it ensures that a one dollar decline in the pre-tax income only translates into a smaller decline in post-tax income, and thus reduces the magnitude of income risk households face. This lower risk translates into a smaller demand for precautionary savings by households, which in turn puts upward pressure on the natural rate of interest.

To quantitatively evaluate the effect of the mechanism described above, the paper uses the model to evaluate the effect of the increase in tax progressivity observed in Australia since the Global Financial Crisis (GFC). It finds that the increase in tax progressivity led to a small increase in the natural rate of interest by only about **0.05 percentage points**. At first glance, it may appear that the model predicts such a small increase in $r^\star$ simply because the underlying increase in tax-progressivity in Australia following the GFC was small. However, this is not entirely accurate. In general, the model predicts very modest increases in $r^\star$ even if one feeds large changes in progressivity into the model. In fact, as Figure 4 of the paper shows, even a large increase in tax progressivity (moving from a proportional tax system to one which features substantial redistribution) would raise r^* by a very modest amount.

It is worth emphasizing that there is large uncertainty surrounding all empirical estimates of $r^\star$. Figure 1, which is taken from Brand et al. (2025), depicts this large uncertainty in the measures of $r^\star$

for the Euro area. For example, as the figure shows, the estimates of $r^\star$ in 2002 lie between -2.5% and 1.5%. Given this uncertainty, an increase in $r^\star$ by a mere 0.05 p.p. makes it hard to argue that monetary policy needs to change its conduct due to changes in the tax progressivity.

2.2 How does tax-progressivity affect the transmission of monetary policy?

The second exercise that the paper undertakes is to see how tax progressivity alters the effect of a monetary tightening on output and inflation. Again, to fix ideas about the forces which shape this interaction, it is useful to use a simple Two Agent New Keynesian (TANK) model based on Bilbiie (2008). Consider an economy in which a fraction λ of households are *hand-to-mouth* (HtM), who consume all their current income, while the remainder are unconstrained households who participate in asset markets. Further, based on empirical evidence, assume that the income of these HtM households is more exposed to fluctuations in aggregate economic activity. In particular, let's assume that a 1% decrease in national income reduces the pre-tax income of this group by $\chi > 1\%$:

$$\frac{d \log y_{\text{htm},t}^{pre-tax}}{dy_t} = \chi > 1,$$

where $y_{\text{htm},t}^{pre-tax}$ denotes the income accruing to HtM households at date t, while y_t denotes national income or GDP.

Since HtM households do not borrow or save, their behavior is not directly affected by changes in monetary policy. Despite this, their presence shapes the aggregate effect of a monetary tightening. Suppose the monetary tightening causes each of the $1-\lambda$ unconstrained households to reduce their spending by 1 dollar, which results in a decline in spending by $1-\lambda$ dollars. Following the standard Keynesian logic, this reduction in spending reduces national income by $1-\lambda$ dollars. However, even though the HtM households are not directly affected by the monetary tightening, the ensuing decline in spending by unconstrained households results in a decline in each HtM household's pre-tax income by $\chi(1-\lambda)$ dollars. Since these households have high marginal propensities to consume, this decline in income results in a one-for-one decline in their spending. This decline in spending further lowers their income and generates a Keynesian multiplier effect. Incorporating all these effects, the effect of the monetary tightening can be written as:

$$dy_t = \frac{1-\lambda}{1-\lambda\chi} dr_t,$$

where dr_t measures the increase in the real interest rate as a result of the monetary tightening.

Importantly, this effect is bigger, the larger is χ.

How does tax progressivity interact with this logic? Progressive taxation weakens this multiplier by compressing the post-tax income distribution, which in the framework above is captured by a smaller after-tax χ.[1] In words, when the tax system is more progressive, HtM households' exposure to aggregate fluctuations declines, so a monetary policy tightening causes output to decline by less.

Qualitatively, the model's predictions line up with the mechanisms I described above. Quantitatively, though, the effects are quite small: in the simulations, greater tax progressivity leads to only negligible differences in output and inflation dynamics. Indeed, as Figure 7 in the paper illustrates, the impulse responses to a monetary tightening look almost identical across the various tax progressivity cases.

These findings suggest that, within this model, monetary policy may be largely unaffected by the degree of tax progressivity. However, as I discuss in the next section, there are reasons to treat this conclusion with some caution.

3. Methodological Reflections

In line with the findings of the paper, my prior is that changes in the progressivity of taxation are of second-order importance for the conduct of monetary policy. That said, I would be cautious about treating the conclusions of this paper as the definitive answer on how tax progressivity interacts with monetary policy.

To be clear, this caution does not stem from any concern about the quality of the paper's analysis. Rather, it reflects a broader unease with the sensitivity of results in the rapidly growing HANK literature to the specific modeling choices researchers make. Relative to the more standard Representative Agent New Keynesian (RANK) framework, HANK models can produce markedly different predictions depending on the way certain features—such as the specification of fiscal policy, the cyclicality of income heterogeneity, or the treatment of risk—are implemented. Even small changes to these features within a given HANK model can materially alter the model's predictions, whereas results from RANK models tend to be more stable under similar modifications.

The sensitivity of HANK models should not be viewed as a flaw. If anything, their prediction that even small changes in fiscal policy can have meaningful macroeconomic responses makes them more realistic than RANK models, where Ricardian equivalence suppresses such channels. However, this same feature demands greater precision and transparency in how fiscal behavior is modeled, as minor

[1] See Alves and Acharya (2024) for a detailed exposition.

inaccuracies can have amplified effects on equilibrium outcomes.

While this may seem like a cheap criticism which can be applied to most modeling exercises, I think that it is particularly important given the recent trend of central banks wanting to add HANK models to the suite of models that they use for policy analysis. The reason I worry about this facet the most is because while HANK models provide a rich depiction of household and firm heterogeneity, the fiscal side of these models is often represented in a highly stylized way. Typically, fiscal policy is summarized by a single reduced-form "fiscal rule" which describes how the government adjusts primary surplus in response to changes in the level of public debt. This paper also follows that common approach, and summarizes the entirety of fiscal policy in terms of a simple fiscal rule.

Although this simplification is standard, it can materially influence the model's predictions. In this paper, for example, footnote 6 notes that replacing the baseline fiscal rule with an alternative specification used by Bayer et al. (2024) leads to instability in some impulse responses, particularly at long horizons where variables no longer converge to steady state. In other words, it seems that depending on the exact specification of the fiscal rule, the model can even predict that monetary policy is not neutral in the long run: a temporary monetary policy shock can permanently move the economy to a different steady state. However, under the baseline specification of the fiscal rule, a temporary monetary policy shock does not have permanent effects. Such divergence in long-run behavior underscores how sensitive the results can be to seemingly small changes in fiscal assumptions and raises legitimate questions about the robustness of the conclusions drawn from any single specification.

To reiterate, this issue is not unique to this paper. The HANK literature has broadly adopted this approach of specifying fiscal policy in this parsimonious fashion. I suspect that this modeling approach is just carrying on from the similar practice in RANK models. However, the failure of Ricardian equivalence in HANK economies implies that their predictions are much more sensitive to the exact fiscal formulation. This greater sensitivity highlights the need for careful modeling and robustness analysis when interpreting results from HANK models.

The need for robustness analysis also raises a broader question. Since a wide range of model ingredients can potentially affect conclusions in important ways, which dimensions should a modeler prioritize when assessing robustness? This is, admittedly, a difficult question. However, it becomes more manageable when one adopts a *problem-driven*, rather than *model-driven*, approach. Starting from a clearly defined policy question helps identify the minimal set of model features necessary to address it credibly. Such an approach clarifies which aspects of heterogeneity matter most for monetary policy and prevents the model from becoming a theoretical black box.

3.1 Welfare Analysis

The paper also analyzes how tax progressivity alters the distributional effects of monetary policy. Specifically, it presents an exercise comparing the welfare impacts of a monetary tightening across households of different income and wealth, under alternative degrees of tax progressivity. The paper finds that the welfare of the poorest households becomes more *sensitive* to monetary shocks, and summarizes this finding as "with higher tax progressivity, the distributional effects of monetary policy are larger." I believe that this conclusion may be somewhat misleading. A greater sensitivity of welfare to shocks does not necessarily imply greater vulnerability or lower welfare overall.

To illustrate, consider a simple example. Suppose that following a monetary policy shock the welfare of the poorest quartile declines by 2 percent more in an economy with more progressive taxation, relative to an economy with less progressive taxation. However, suppose that the higher degree of tax progressivity also raises the welfare of the poorest quartile by 5% on average because of increased redistribution. If we focus only on sensitivity, we would conclude that poor households are more exposed, but this ignores the fact that their overall welfare is higher in every state of the world. Thus, a more progressive tax system can make welfare more responsive to shocks while simultaneously improving average welfare levels. This demonstrates that one cannot obtain an accurate picture of whether the distributional effects of monetary policy are larger under greater progressivity by looking only at the sensitivity of welfare while ignoring differences in average welfare levels.

This issue reflects a deeper methodological limitation. As Benigno and Woodford (2005) show, comparing welfare across policies that produce different average outcomes requires dynamics accurate to at least second-order. Unfortunately, most current solution methods used to solve HANK models (such as the Reiter approach or the sequence-space Jacobian technique) produce first-order accurate dynamics by construction.[2] The methods used in this paper also only produce first-order-accurate dynamics of aggregate macroeconomic variables and thus, cannot be used to get accurate estimates of welfare effects of changes in tax progressivity, which take into account the effects of progressivity on both the average level and sensitivity of welfare.

4. Conclusion

In summary, HANK models can be very useful tools for central banks trying to fully understand the transmission mechanism of monetary policy and its redistributive implications. But as I mentioned

[2]While there are some attempts to develop methodology which can be used to solve HANK models to generate at least second-order-accurate dynamics (see, e.g., Bhandari et al. (2023)), such methods have not been adopted as much as the ones described above.

earlier, I think that there are serious methodological challenges which need to be addressed before policymakers can be fully confident that the HANK model they are using identifies the key policy trade-offs correctly. Complex models can generate insights into the distributional channels of monetary policy, but they also risk becoming detached from the practical needs of policymakers. Going forward, it will be important to build frameworks that balance realism with transparency and theoretical discipline.

Focusing on the topic of this paper, progressive taxation remains a key tool for achieving distributional objectives. But as far as the conduct of monetary policy is concerned, its effects are likely second-order. For the Reserve Bank of Australia—and for central banks more generally—the main lesson is that monetary and fiscal policy interact, but not all interactions are quantitatively important. Recognizing which ones matter most is the essence of sound macroeconomic analysis.

References

Alves, Felipe and Sushant Acharya, "How changes in the share of constrained households affect the effectiveness of monetary policy," Technical Report, Bank of Canada 2024.

Bayer, Christian, Benjamin Born, and Ralph Luetticke, "Shocks, frictions, and inequality in US business cycles," *American Economic Review*, 2024, *114* (5), 1211–1247.

Benigno, Pierpaolo and Michael Woodford, "Inflation Stabilization and Welfare: The Case of a Distorted Steady State," *Journal of the European Economic Association*, 2005, *3* (6), 1185–1236.

Bhandari, Anmol, Thomas Bourany, David Evans, and Mikhail Golosov, "A perturbational approach for approximating heterogeneous agent models," Technical Report, National Bureau of Economic Research 2023.

Bilbiie, Florin O, "Limited asset markets participation, monetary policy and (inverted) aggregate demand logic," *Journal of economic theory*, 2008, *140* (1), 162–196.

Brand, Claus, Noëmie Lisack, and Falk Mazelis, "Natural rate estimates for the euro area: insights, uncertainties and shortcomings," *Economic Bulletin Boxes*, 2025, *1.*