HANK: A New Core of Usable Macroeconomics

Adrien Auclert*

January 2025

Abstract

I argue that Heterogeneous-Agent New Keynesian ("HANK") models are a natural contender to define a new core of usable macroeconomics, replacing the ISLM/Mundell-Fleming paradigm prevailing after the 1960s and the New Keynesian paradigm prevailing after the 1990s. HANK is "macro" (general equilibrium, aggregating up from behavior of optimizing units), it is "usable" (featuring rational expectations, consistent with the known macro effects of monetary and fiscal policy, and the micro behavior of consumption), and it has a "core" version (stripped down, easy to teach and build upon). I discuss what we know about this new core, what we are starting to learn, and where more research is still needed.

What makes up a core of usable macroeconomics? The answer to this question has changed over time, but I think macroeconomists of all generations would agree with the following definitions. To qualify as *macroeconomics*, it surely has to be about general equilibrium, with many markets interacting. To be actually *usable* by policymakers, it has to be consistent with what we know (or think we know) about the effects of monetary and fiscal policy. And, to meet the *core* requirement, it has to be a sufficiently stripped-down model, one that is easy to teach and build upon.

For decades, the model that best satisfied these three criteria was the ISLM model, or its open-economy extension, the Mundell-Fleming model. Members of the 1997 AEA panel that were asked the same question agreed as much. This model was widely taught and ubiquitous in policymaking and in economics journals from the 1950s to the 1980s. Its influence has been so long-lasting that it is still taught to undergraduate macro students and still used by some policymakers today, especially in fiscal policy circles.

The rational expectations revolution that started in the late 1970s led to a dramatic shift in the definition of the core of usable macro. By macro, came to be meant not just an equilibrium model, but one based on the behavior of optimizing agents. By usable, came to be meant a model that

*Stanford University and NBER. Email: aauclert@stanford.edu.

Prepared for the AEA Papers and Proceedings 2025. I thank Aysegul Sahin for organizing the session, as well as Marty Eichenbaum, Bart Hobijn, Matt Rognlie and Ludwig Straub for helpful comments

was built on rational expectations, so that agents in the model couldn't continuously be fooled by policymakers. And by core, came to be meant some version of the neoclassical growth model.

Clearly these new criteria implied abandoning the ISLM model for an alternative core. On the basis of the new criteria alone, the real business cycle (RBC) model would probably have fit the bill. But that model didn't fit the *earlier* usability criterion, since it was not consistent with the real effects of monetary policy in the data. Augmenting the RBC model with nominal rigidities was the most natural way to achieve that. And so, the model that replaced ISLM as the core of usable macroeconomics starting in the 1990s was the New Keynesian (NK) model, with its three-equation form taught in most PhD courses and summarized in classic textbooks (eg. Jordi Galí 2008).

This was more or less the state of affairs when the 1997 AEA panel concluded. But since then, another big revolution has taken place. Large-scale administrative datasets have become increasingly available, showing a vast amount of heterogeneity at the micro level, and with this micro-data revolution came important developments in applied microeconomics that allowed for credible estimation of causal relationships.

It became increasingly desirable for our models to be consistent with the heterogeneity in the micro data: in addition to the criteria above, *macro* came to mean aggregating up from the behavior of heterogeneous units. It also became desirable for our models to match these new, credibly estimated causal moments. Particularly relevant for macro was the literature estimating marginal propensities to consume out of transitory income shocks, or MPCs for short (eg, David S. Johnson, Jonathan A. Parker and Nicholas S. Souleles 2006). This literature documented very large MPCs that clearly rejected the representative-agent assumption underpinning the New Keynesian model, with its implied permanent-income behavior and extremely low MPCs. A *usable* model had to be a model consistent with the MPCs in the data. Luckily, by that point there was an off-the-shelf core model consistent with high MPCs (at least under some calibrations): the canonical heterogeneous-agent (HA) model that had been developed by S. Rao Aiyagari (1994) and Per Krusell and Anthony A. Smith (1998), among others.

But there again, the earlier usability criterion was binding. Nominal rigidities were still the simplest known way to make the model consistent with the effects of monetary policy in the data. This is where HANK models came in, as a synthesis between the HA paradigm consistent with the heterogeneity and MPCs in the data, and the NK paradigm that was usable for monetary and fiscal policy analysis.

1 The need for a new paradigm

Two additional developments of the past few decades contributed to the emergence of HANK models as a new paradigm.

Rethinking fiscal policy. The first was the increasing recognition that fiscal policy—and fiscal transfers in particular—had a core role to play in aggregate macroeconomic stabilization. Conventional wisdom at the time of the 1997 AEA panel was that monetary policy should play the

main stabilization role, with fiscal policy relegated to other objectives such as distribution. But when monetary policy hit the zero lower bound on interest rates, first in Japan in the late 1990s, and then in many advanced economies after the global financial crisis a decade later, it became clear that it had to be supplemented with other aggregate stabilization tools. Fiscal transfers to households became a key tool for preventing economic collapse, and policymakers have used them with increasing frequency—most recently in the wake of the Covid recession in 2020-2021, and again in Europe during the energy price crisis of 2021-2022.

The general takeaway from these episodes was that fiscal transfers did not just boost economic activity and inflation as intended, but that they did so in a persistent (and perhaps unintended) way. There is ongoing research establishing these facts more formally, but I think the consensus is well captured by a newspaper article from July 2023, two full years after the end of the pandemic stimulus transfers. Commenting on the resilience of both consumer spending and inflation in the face of the Fed tightening, the paper argued: "this resilience is partly a result of government stimulus programs during the COVID-19 pandemic. These programs generated a stockpile of excess savings that has continued to support household spending".[1]

These effects cannot be explained by either of the two earlier core models of fiscal policy. The New Keynesian model is a non-sequitur for studying the macroeconomic effects of deficit-financed fiscal transfers: since Ricardian equivalence holds in that model, households save any fiscal transfer in anticipation of future taxes, so that private savings increase with no consequence for spending. But the ISLM model is also inappropriate, since it is completely static: in this model, there is a nontrivial multiplier for deficit-financed transfers, but (at least in modern formal versions of this model with hand-to-mouth consumers) the effects on economic activity and inflation disappear right after the transfers do.

This created space for a new paradigm for fiscal policy: one in which, consistent with the empirical evidence, fiscal transfers are partially spent and partially saved, with households later spending down their "excess savings", leading to persistent effects on economic activity and inflation. As I will argue below, this is exactly what HANK delivers.

Rethinking monetary policy. The other major development was the increasing recognition that monetary policy had distributional effects that policymakers needed to openly discuss. While I think that these effects were always understood to be present to some extent, the dominant view had been that these were small side effects that were not part of the mission of monetary policymakers, and therefore were best left out of our models.

However, the increasing focus on inequality in the public discourse, exemplified by the success of Thomas Piketty's book (Thomas Piketty 2014), created pressure for policymakers to address these effects publicly, which in turn raised the demand for models in which these effects were explicitly acknowledged and could be studied quantitatively.

And this is indeed what central bankers have done, as narrated, for instance, by ECB executive board member Isabel Schnabel: "we analyze the interaction between income & wealth distribu-

[1] "Pandemic Stimulus Still Supporting Spending", USA Today, 11 July 2023.

tions and monetary policy as part of the policy process. Distributional effects matter for monetary policy transmission and are part of our proportionality analysis".[2] This type of analysis is exactly what HANK is designed for.

2 HANK as a new core model

These factors all point to the promise of HANK as a new paradigm. Meanwhile, objective measures of adoption are consistent with it taking over as a (the?) new core model of macro. Today, HANK is taught in some form in second-year PhD courses at many universities. The paper that gave its literature its name, Greg Kaplan, Benjamin Moll and Giovanni L. Violante (2018), has nearly 2000 citations on Google Scholar. Many central banks and a number of fiscal policy institutions have a HANK development team. And numerical toolboxes are increasingly available that automate key solution steps: in addition to the sequence-space Jacobian toolbox that my coauthors and I developed (Adrien Auclert, Bence Bardóczy, Matthew Rognlie and Ludwig Straub 2021), there is the BASE for HANK toolbox (Christian Bayer, Benjamin Born and Ralph Luetticke 2024), and a couple of other projects under development. Obviously, these solution toolboxes are a key part of widespread usability.

The RBC model really became a core model of macroeconomics when subsequent efforts were made to strip down the original model presented by Finn E. Kydland and Edward C. Prescott (1982) to its most essential ingredients. Similarly, trial and error has led the HANK literature to deviate from the formulations proposed by early papers in the literature, such as Kaplan, Moll and Violante (2018), in favor of a simpler version that contains only the essence. In a recent review article, my coauthors and I propose a "canonical" HANK model that we argue can serve as this relevant core (see Adrien Auclert, Matthew Rognlie and Ludwig Straub 2025).[3] A distinctive feature of the literature is that the models are often presented and analyzed in what we call the *sequence space*, where perfect-foresight transition paths are studied after aggregate shocks. Due to certainty equivalence, there is no loss of generality to doing this when aggregate shocks are small, and the payoff is an enormous gain in terms of intuition for economic mechanisms and in simplifying the solution algorithms.

In the simplest possible version of this model, agents face uninsured idiosyncratic earnings risk captured by a time-varying level of skill e_{it} and a borrowing constraint $\underline{a}$. Each period, given bond holdings a_{it-1} and hours of work n_{it}, they receive post-tax labor income $(1-\tau_t)\, w_t e_{it} n_{it}$ and interest income $(1+r_{t-1})\, a_{it-1}$, and choose between consumption and savings for the next period to maximize $\mathbb{E}_0 \sum_t \beta^t \{u(c_{it}) - v(n_{it})\}$ by choice of c_{it}, a_{it} subject to the borrowing constraint $a_{it} \geq$

[2] Post on X (formerly Twitter), 8 January 2024.

[3] An important, parallel development has been that of tractable versions of the basic HANK framework, which make simplifications to remove the wealth distribution as a state variable but retain some of its core economic mechanisms, such as the macroeconomic importance of precautionary savings. See, eg, Florin O. Bilbiie (2024).

$\underline{a}$ and the budget constraint

$$c_{it} + a_{it} = (1 + r_{t-1})\, a_{it-1} + (1 - \tau_t)\, w_t e_{it} n_{it}$$

The government sets the tax rate τ_t, spends G_t on goods and services, and issues bonds B_t to satisfy a period budget constraint $B_t + \tau_t w_t \int e_{it} n_{it} di = (1 + r_{t-1})\, B_{t-1} + G_t$. Production is from labor, $Y_t = N_t$, and there are flexible prices and perfect competition, so that the pre-tax wage per unit of skill is $w_t = 1$. Finally, goods and asset markets clear, so that $\int c_{it} di + G_t = Y_t$ and $\int a_{it} di = B_t$.

So far, this is a completely standard general equilibrium model in the classic heterogeneous-agent tradition. To make the model usable, let us now introduce nominal wage rigidities: accordingly, labor is rationed to be $n_{it} = N_t$ for all individuals i. Staggered nominal wage setting then implies a standard New Keynesian Phillips curve for nominal wage inflation, and therefore also for price inflation (recall that here the real wage is constant). Equilibrium given an exogenous specification of fiscal policy $\{G_t, B_t\}$ and monetary policy $\{r_t\}$ is a set of aggregates, optimal decisions, and evolution of the distribution of agents consistent with these decisions, such that all markets clear at all dates.

A fascinating aspect of this model is that it behaves very similarly, in many ways, to the old ISLM model, but with some important differences that result from its microfoundations. Consider for instance the effect of fiscal policy, which we will isolate by assuming that monetary policy keeps a constant $r_t = r$ at all dates. Given our assumptions, labor income for individual i at t is $(Y_t - T_t)\, e_{it}$, where $T_t = (1 + r)\, B_{t-1} + G_t - B_t$ denotes the (exogenous) amount of aggregate tax income, and so the only aggregate variable that enters the household problem is $Y_t - T_t$. But that implies that aggregate consumption at time t, C_t, only depends on the sequence of aggregate post-tax income $Y_s - T_s$ at all dates $s = 0, \ldots, \infty$, via a function we will call $\mathcal{C}_t$. And so, equilibrium is simply a fixed point in the sequence space: we have to find the sequence $\{Y_t\}$ that, given exogenous $\{G_t, T_t\}$, satisfies at all dates

$$\mathcal{C}_t\left(\{Y_s - T_s\}\right) + G_t = Y_t$$

This looks very similar to the standard Keynesian cross from the ISLM model, except that is is "intertemporal" (Adrien Auclert, Matthew Rognlie and Ludwig Straub 2024). The basic intuition is the same: increases in government spending or reductions in taxes raise aggregate demand and therefore income, feeding back into spending through a multiplier process. But because households optimize spending over time, now past income as well as anticipation of future income matters for current spending. In fact, in this framework, past deficit-financed tax cuts can boost spending for many years, capturing the prolonged impact of excess savings on activity and inflation that was lacking from earlier generations of models, with the evolution of the wealth distribution being central to this process (see Adrien Auclert, Matthew Rognlie and Ludwig Straub 2023). To quantify these effects, the model can be calibrated so that its *intertemporal MPCs* (eg $\frac{\partial \mathcal{C}_t}{\partial Y_0}$) are

consistent with the large, credibility-revolution-robust estimates that we have from the microdata.

3 Current frontiers and future directions

I have argued that HANK is firmly establishing itself as a new core model for analysis of monetary and fiscal policy. Just like the New Keynesian framework, this core is tractable and easy to augment to study a wide variety of topics.

The core model is qualitatively consistent with the known effects of monetary policy (say, the aggregate effects on output and inflation of shocks to r_t), though for well-understood reasons it falls short of capturing the quantitative dynamic patterns. But medium-scale versions of these models are in development, both in academic research and in HANK development teams at policy institutions, that remedy this shortcoming. The core model features rational expectations, but it can also be augmented to deal with a large class of departures from rational expectations that have been documented in the data. The core model is one of a closed economy, but it can be augmented with a demand for imports and exports from the rest of the world, so that it can speak to the impact of monetary policy on exchange rates, of exchange rate or world energy price fluctuations on the real income of consumers, or of the effects of fiscal policy on the balance of trade. The core model does not feature any aggregate uncertainty, but it can be augmented to consider portfolio choice between multiple assets, as well as the risk-based pricing of assets. The core model is positive, but we can exploit the welfare objective from within the model to derive lessons for optimal policy.

These are all places where the literature has already made significant progress. By contrast, less progress has been made on a few essential issues. The description of the labor market remains quite limited in the baseline model, even when it is augmented with a search-and-matching module. The wage Phillips curve is an extremely unsatisfactory part of the model: it does a poor job at explaining both aggregate wage inflation and the heterogeneity in wage adjustment after shocks in the data. The average MPC level is consistent with the data, but the model lacks reasons for why MPCs may be high that aren't related to liquidity—overstating the empirical relationship between MPC and low liquidity. Finally, the main lessons for optimal policy in this framework are essentially still unknown, with important barriers to establishing core results such as the lack of existence of a Ramsey steady-state in baseline specifications.

These are all exciting frontiers to work on, which perhaps will define the next version of a core of usable macro.

References

Aiyagari, S. Rao. 1994. "Uninsured Idiosyncratic Risk and Aggregate Saving." *Quarterly Journal of Economics*, 109(3): 659–684.

Auclert, Adrien, Bence Bardóczy, Matthew Rognlie, and Ludwig Straub. 2021. "Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models." *Econometrica*, 89(5): 2375–2408.

Auclert, Adrien, Matthew Rognlie, and Ludwig Straub. 2023. "The Trickling Up of Excess Savings." *AEA Papers and Proceedings*, 113: 70–75.

Auclert, Adrien, Matthew Rognlie, and Ludwig Straub. 2024. "The Intertemporal Keynesian Cross." *Journal of Political Economy*, 132(12): 4068–4121.

Auclert, Adrien, Matthew Rognlie, and Ludwig Straub. 2025. "Fiscal and Monetary Policy with Heterogeneous Agents." *Annual Review of Economics*, forthcoming.

Bayer, Christian, Benjamin Born, and Ralph Luetticke. 2024. "Shocks, Frictions, and Inequality in US Business Cycles." *American Economic Review*, 114(5): 1211–1247.

Bilbiie, Florin O. 2024. "Monetary Policy and Heterogeneity: An Analytical Framework." *Review of Economic Studies*, forthcoming.

Galí, Jordi. 2008. *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework.* Princeton University Press.

Johnson, David S., Jonathan A. Parker, and Nicholas S. Souleles. 2006. "Household Expenditure and the Income Tax Rebates of 2001." *American Economic Review*, 96(5): 1589–1610.

Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante. 2018. "Monetary Policy According to HANK." *American Economic Review*, 108(3): 697–743.

Krusell, Per, and Anthony A. Smith. 1998. "Income and Wealth Heterogeneity in the Macroeconomy." *Journal of Political Economy*, 106(5): 867–896.

Kydland, Finn E., and Edward C. Prescott. 1982. "Time to Build and Aggregate Fluctuations." *Econometrica*, 50(6): 1345–1370.

Piketty, Thomas. 2014. *Capital in the Twenty-First Century.* . 1st ed., Cambridge, MA:Belknap Press.