Research Statement

Raphaël Huleux

September 2026

I am a macroeconomist working at the intersection of inequality, stabilization policy, and computational economics. My research uses quantitative heterogeneous-agent models to study how micro-level heterogeneity shapes the wealth distribution, asset prices, and the transmission of fiscal and monetary policy. Methodologically, I combine sequence-space Jacobian (SSJ) techniques with GPU-accelerated deep-learning methods to solve large, non-linear dynamic problems that are otherwise computationally intractable. My work follows two lines. The first asks how the rise in permanent labor income inequality changes macroeconomic outcomes and policy. The second develops deep-learning solution methods for heterogeneous-agent models, and uses them to ask when simpler methods are good enough.

Inequality and the transmission of stabilization policies (JMP)

In *Permanent Labor Income Inequality and Stabilization Policies*, with Eustache Elina (Banque de France), we study how the distribution of permanent labor income (PLI) affects the transmission of fiscal and monetary policy. In standard heterogeneous-agent models, hand-to-mouth status is unrelated to PLI, and the PLI distribution has no effect on aggregate dynamics. We show that this is at odds with the data: in the PSID, the hand-to-mouth share falls from 64% to 22% across PLI quintiles, and this gradient has steepened since 1999.

We match this gradient in a heterogeneous-agent New Keynesian model with a non-homothetic taste for wealth. High-PLI households save more because of the taste for wealth, and their savings push down the equilibrium interest rate. For low-PLI households, the taste for wealth is negligible and they behave as buffer-stock savers. Facing a lower interest rate, they save less and are more often at the borrowing constraint. The model therefore generates a negative correlation between PLI and MPCs. The MPC out of labor income, which weights households by their income share, is lower than the MPC out of a lump-sum transfer, which weights them equally. A dollar of labor income generates 22% less spending in present value than a dollar of transfers, against 3% in a standard model.

We have two main findings. First, matching the gradient makes the incidence of policy a first-order determinant of the cumulative multiplier: the same deficit can stimulate or contract output depending on who receives it and who repays it. The same holds for monetary policy, since a rate cut lowers the interest bill on public debt and its effect on output depends on how the government rebates these savings. Second, the rise in PLI inequality since 1970 has amplified these effects. The cumulative multiplier of a transfer repaid by taxes rose from 0.82 to 0.93, the multiplier of a tax cut repaid by transfer cuts fell from -0.49 to -0.61, and the output response to a monetary shock became about 20% more sensitive to the fiscal rule. Higher inequality therefore increases the need for coordination between fiscal and monetary policy.

From income to wealth inequality

In *From Income to Wealth Inequality: Trickle-Down vs. Capital Gains* (submitted), with Eustache Elina, we study how higher PLI inequality feeds back on the wealth distribution through asset prices. Households with higher permanent income save a larger share of it, so a more unequal PLI distribution raises aggregate saving, and prices must adjust for asset markets to clear. If the extra saving finances capital, interest rates fall, wages rise, and poorer households gain. If it raises asset prices instead, the capital gains accrue to the rich.

In a heterogeneous-agent model with non-homothetic wealth preferences and imperfect competition, markups weaken the first channel and let the second dominate for five decades along the transition. Fed with the observed U.S. rise in inequality since 1970, the model generates two fifths of the rise in the wealth-to-output ratio and a quarter of the rise in Tobin's Q. These price effects raise the average wealth of the top 0.1% by 25% and average wealth by 8% by 2020, while under perfect competition both fall. The paper clarifies when "more saving by the rich" trickles down through physical capital and when it mostly capitalizes into asset prices.

Deep learning for large life-cycle models

In *Deep Learning Solutions of Large Non-Convex Life-Cycle Models*, with Jeppe Druedahl and Jacob Røpke, we introduce a deep-learning algorithm for life-cycle models with both continuous and discrete choices. We work directly on the Bellman equation, approximate both value and policy functions with neural networks, and train on a simulated sample instead of tensor-product grids. We test the algorithm in a consumption-saving model with multiple durable goods and non-convex adjustment costs, which can still be solved with standard dynamic programming. The comparison shows how the algorithm alleviates the curse of dimensionality: adding one durable good raises the solution time of dynamic programming from under 20 minutes to over 13 hours, while our algorithm needs only 40 more minutes to reach the same accuracy. We then solve a large life-cycle model with labor supply and human capital accumulation, a portfolio choice between a risky and a risk-free asset, and housing and mortgage choices, in 12 hours on a single GPU.

Global solutions of HANK models with aggregate risk

In *Fiscal Multipliers in Globally Solved HANK Models with Aggregate Risk*, with Jeppe Druedahl and Jacob Røpke, we ask how much expectations of aggregate risk matter for the dynamics of HANK models. These models are usually linearized, which rules out any role for aggregate risk and any dependence of fiscal multipliers on the sign of the shock, its size, or the state of the economy. We compute the global solution of two HANK models with deep learning, writing policies and aggregate outcomes as functions of the history of aggregate shocks, and validate its accuracy against the non-linear SSJ solution. Our main result is that the non-linear perfect-foresight solution, computed with standard fast methods, closely approximates the global solution, except for a level shift driven by precautionary saving. This holds in calibrations where the effects of fiscal policy depend strongly on the sign, size, and state. For applied work, this means the standard perfect-foresight workflow is sufficient to study fiscal policy in this class of models.