Research Statement

Andreas Schaab

I am a macroeconomist working on (1) the aggregate and normative implications of household heterogeneity and (2) monetary policy. Much of my recent work addresses the question how we can make aggregate welfare assessments and design optimal policy in environments with heterogeneous individuals.

Normative macroeconomics with heterogeneous agents. Identifying the sources of welfare gains and losses is critical to assess the impact of shocks and the desirability of policy interventions. In environments with a single individual, such welfare assessments are straightforward. In the presence of rich heterogeneity, however, this becomes a challenging task. In particular, assessing the aggregate welfare consequences of policies that are not Pareto-improving requires interpersonal welfare comparisons, trading off gains and losses across different individuals. While the typical approach of making such comparisons using social welfare functions (SWFs) is widely applicable, it is difficult to identify how such a criterion makes tradeoffs across heterogeneous individuals.

In "Welfare Assessments with Heterogeneous Individuals" [4] (R&R at JPE), Eduardo Dávila and I develop a new approach to make aggregate welfare assessments based on dynamic stochastic welfare weights. Our main result additively decomposes any aggregate welfare assessment into four distinct normative considerations: aggregate efficiency, risk-sharing, intertemporal-sharing, and redistribution. This decomposition applies to general dynamic stochastic environments with heterogeneous individuals. The paper characterizes a host of desirable properties satisfied by our decomposition. For example, we show that (i) in complete markets economies, welfare gains and losses can only emerge from aggregate efficiency and redistribution; (ii) in economies with riskfree borrowing and lending, welfare gains can never emerge from intertemporal-sharing; (iii) different SWFs exclusively disagree on redistribution; (iv) Pareto improvements must increase aggregate efficiency, risk-sharing, or intertemporal-sharing; and (v) the aggregate efficiency, risk-sharing, and intertemporal-sharing components are invariant to monotonically increasing transformations of individuals' instantaneous utilities.

Identifying the ultimate origins of welfare gains and losses is further complicated in the presence of rich production networks that feature pure intermediates. And while the well-established growth accounting approach seeks to trace the origins of growth in final output, welfare and output do not coincide. In "Welfare Accounting" [6] (with Dávila), we develop a new decomposition of welfare assessments that applies to general economies with heterogeneous individuals and disaggregated production. In contrast to existing work, our welfare accounting decomposition is exclusively based on preferences and technologies; it does not rely on assumptions about the (optimizing) behavior of agents, firm objectives, individual budget constraints, prices, or the notion of equilibrium. A central contribution of this decomposition is to identify a small set of welfare-relevant statistics that allow us to map arbitrary changes in allocations to welfare: marginal rates of substitution (MRS), marginal welfare products (MWP), and the marginal social value of output (MSV). We then leverage our decomposition to characterize efficiency in our environment, providing, to our knowledge, the first general characterization of efficiency conditions for disaggregated production economies with heterogeneous individuals. Finally, we specialize welfare accounting to competitive economies with wedges, which allows us to map the welfare-relevant

statistics we have identified to measurable prices (and wedges). We provide a new general Hulten's theorem for environments with heterogeneous individuals, elastic and fixed factors, and arbitrary social welfare functions. Its generality allows us to systematically discuss the many qualifications associated with this result. In particular, we show that Hulten's theorem applies to frictionless competitive economies, rather than efficient economies as is often stated. Similarly, Hulten's theorem is fundamentally a result about aggregate efficiency, rather than final output or welfare. Finally, we develop a new converse result to Hulten's theorem that establishes the necessary condition under which the marginal social value of a good is given by its market price.

In ongoing and future work, I hope to advance this agenda on normative macroeconomics with heterogeneous agents in several directions. In [8] (with Sergi Barcons and Dávila), we study intergenerational welfare assessments and extend the framework developed in [4] to overlapping generations environments. In [9] (with Dávila), we extend the welfare accounting approach of [6] to dynamic stochastic environments. Most importantly, I hope to leverage this framework to revisit the welfare implications of cross-sectional heterogeneity in applied and policy-relevant contexts. An important application is monetary policy and the design of central bank mandates.

Monetary policy and optimal policy design. Just prior to the Covid-19 pandemic, both the Federal Reserve and the European Central Bank engaged in multi-year, long-term strategic reviews. These reviews were in part motivated by the persistent decline in the natural rate of interest, a seemingly flattening Phillips curve, and heightened appreciation for the distributional consequences of monetary policy. The Federal Reserve concluded its review in August 2020, switching to an average inflation target and updating the definition of its employment mandate as a "broad-based and inclusive goal."

In "Optimal Monetary Policy with Heterogeneous Agents" [3] (R&R at AER, with Dávila), characterize optimal monetary policy in a canonical heterogeneous-agent New Keynesian (HANK) model under a utilitarian SWF. Our approach allows us to systematically revisit the canonical New Keynesian consensus on optimal monetary policy, studying policy under discretion, with commitment, and from a timeless perspective. We show that a utilitarian planner in a HANK economy has an incentive to raise output above its natural level and overheat the economy, even in the absence of markup distortions. This occurs because the planner values redistribution toward indebted, high marginal utility households via lower interest rates. Optimal monetary policy under discretion trades off this novel redistribution motive against aggregate stabilization. And when agents anticipate the planner's incentives to lower interest rates, the redistribution motive becomes a new source of inflationary bias in the sense of Barro and Gordon (1983). We study optimal policy under commitment in three steps. In a first step, we show that it is optimal to commit to 0 inflation in the long run, as in the representative-agent benchmark. In a second step, while the standard Ramsey problem features no inflationary bias in the long run, optimal monetary policy in the short run still overheats the economy due to a time-0 problem. In our setting, time inconsistency in the sense of Kydland and Prescott (1977) arises due to not only the Phillips curve but also a distribution of forward-looking household Euler equations. We extend the Marcet and Marimon (2019) recursive multiplier approach to continuous-time heterogeneous-agent economies and show that a timeless planning solution requires two sets of penalties, the standard inflation penalty and a novel distributional penalty. In a third and final step, we study optimal monetary stabilization policy under the timeless Ramsey problem and show that Divine Coincidence fails even in

the absence of cost-push shocks since the planner perceives a tradeoff between aggregate stabilization and distributional considerations. Lastly, the paper also makes a methodological contribution, extending the sequence-space approach of Auclert et al. (2021) to Ramsey problems and welfare analysis.

In future work, I hope to study optimal monetary policy with distributional considerations in quantitatively more realistic environments. In "Monetary and Fiscal Policy According to HANK-IO" [7], Stacy Tan and I take a first step in this direction, developing a new multi-sector heterogeneous-agent New Keynesian model with an input-output production network. We document systematic household-sector linkages in micro data and calibrate our model to match them. To identify when these linkages have implications for policy transmission, we analytically characterize an as-if benchmark that features a strict decoupling between household and sectoral heterogeneity. Away from this benchmark, novel earnings and expenditure heterogeneity channels emerge that govern the propagation of demand and supply shocks. Quantitatively, these channels shape the transmission of stabilization policy to aggregates, as well as its distributional and sectoral consequences. In [10] (with Dávila), we build on [3] and [4] and characterize the optimal design of central bank mandates with distributional considerations.

In "A Theory of Dynamic Inflation Targets" [2] (R&R at AER), Christopher Clayton and I take a dynamic mechanism design approach to study the question how inflation targets should be designed in the presence of persistent shocks. Historically, many central banks' inflation targets have evolved substantially since their inception in the early 1990s. Most recently, the Fed adopted an average inflation target in 2020 and the ECB switched from an asymmetric to a symmetric target in 2021; both central banks justified these target adjustments as motivated in part by the decline in the natural rate. In academic discourse, an important motivation for inflation targets is the interaction between a time consistency problem and central bank private information: commitment to a rule corrects inflationary bias while flexibility to set inflation allows the central bank to respond to private information about shocks. Prior work has studied static inflation target mechanisms. In order to speak to the empirical regularity that central banks regularly update their targets, our paper develops a theory of dynamic inflation targets in an environment with persistent shocks and persistent private information, motivated by the recent debate on persistent changes in r* and the slope of the Phillips curve. Our main result is that a time-varying dynamic inflation target mechanism implements the efficient, full-information commitment allocation. In our leading applications, we then show that a declining natural rate and a flattening Phillips curve have exactly opposite implications for the optimal response of a dynamic inflation target, calling respectively for an increase and a decrease in target level and flexibility.

After the Great Recession, central banks were confronted with another set of important questions as they took on an increasingly important role in financial stability and oversight: What was the optimal design of bailout and bail-in regimes? What were the interaction between monetary policy and financial regulation? And how should financial stability regulation be extended and coordinated across borders? In earlier work, Christopher Clayton and I studied some of these questions. In "Multinational Banks and Financial Stability" [1] (QJE, 2022), we developed a theoretical framework to study international bank regulation and the scope for international cooperation in macroprudential policies. The paper's main result shows that non-cooperative national governments using revenue-generating Pigouvian taxation can achieve the global optimum. By contrast, setting quantity regulations non-cooperatively fails to achieve the globally efficient outcome, under-regulating domestic banks, and over-regulating foreign banks. In "Bail-Ins, Optimal Regulation, and Crisis Resolution" [5] (R&R at RFS), we develop a dynamic

contracting framework to study bank bail-in regimes. In the presence of a repeated monitoring problem, the optimal bank capital structure combines standard debt, which induces liquidation and provides strong incentives, and bail-in debt, which restores solvency but provides weaker incentives. Optimal policy increases use of bail-in debt when there are fire sales.

Published papers

1. "*Multinational Banks and Financial Stability*," with Christopher Clayton. August 2022.
   Quarterly Journal of Economics, 137(3): 1681-1736.

Under revision

2. "*A Theory of Dynamic Inflation Targets*," with Christopher Clayton. April 2023.
   Revise and resubmit at **American Economic Review**.
3. "*Optimal Monetary Policy with Heterogeneous Agents: Discretion, Commitment, and Timeless Policy*," with Eduardo Dávila. February 2023.
   Revise and resubmit at **American Economic Review**.
4. "*Welfare Assessments with Heterogeneous Individuals*," with Eduardo Dávila. September 2022.
   Revise and resubmit at **Journal of Political Economy**.
5. "*Bail-Ins, Optimal Regulation, and Crisis Resolution*," with Christopher Clayton. August 2022.
   Revise and resubmit at **Review of Financial Studies**.

Working papers

6. "*Welfare Accounting*," with Eduardo Dávila. September 2023.
7. "*Monetary and Fiscal Policy According to HANK-IO*," with Stacy Yingqi Tan. June 2023.

Work in progress

8. "*Intergenerational Welfare Assessments*," with Sergi Barcons and Eduardo Dávila.
9. "*Dynamic Stochastic Welfare Accounting*," with Eduardo Dávila.
10. "*Central Bank Mandates with Distributional Considerations*," with Eduardo Dávila.