Centre for Research and Teaching in Economics (CIDE)

Master's in Economics | Fall 2026

Computational Methods for Macroeconomics

Solving and Estimating DSGE Models

Instructor: Vasudeva Ramaswamy

Office hours: TBD

Format: 14 weeks, two classes per week (28 classes)

Software: Dynare (MATLAB or GNU Octave) and Helix (browser-based, with a Julia package)

Course objective

This course provides an overview of solution and estimation techniques for dynamic stochastic general equilibrium (DSGE) models. We will cover the basics of numerical approximation techniques and introduce latest advances in the field. By the end, students should be able to construct a model, implement the mathematics required to solve it, estimate it using data, and extend it from the representative household to settings with heterogeneous agents.

The course has four parts. Part I develops the mathematical foundations on the whiteboard: deterministic and stochastic difference equations, linear rational-expectations systems and the conditions under which they can be solved (Blanchard–Kahn and related), linearization and the state-space representation, higher-order perturbation, a (very brief) overview of projection and global methods, and sequence-space techniques. Part II turns to model building: a brief treatment of dynamic programming, then the real business cycle and New Keynesian models — derived and solved by hand first, then in Dynare and Helix — and models of permanent shocks and long-run transitions. Part III focuses on empirical implementation: the state-space form and the Kalman filter, Bayesian inference and posterior simulation, and limited-information approaches such as impulse-response matching. Part IV introduces incomplete-markets and HANK models, solved using the sequence-space methods.

The course emphasizes practical hands-on experience. Part I is deliberately kept to the whiteboard — pure mathematics, with analytical problem sets. From Part II onward, most sessions combine conceptual material with computational exercises, and students are expected to build and solve models themselves rather than just run provided code. By the end of the course, students should be able to take an unfamiliar published DSGE model, reproduce it, estimate it, and assess its strengths and limitations. These skills are directly relevant for work in central banks, ministries, and research institutions, and they provide solid preparation for applications to PhD programs.

Prerequisites

This course is aimed at master's students intending to pursue an economics PhD or to work at central banks, ministries of finance, or comparable policy institutions. It assumes prior exposure to dynamic macroeconomics, in particular:

- Macroeconomics I: general-equilibrium models with microfoundations, the Euler equation, infinite-horizon optimization.
- Dynamic Methods for Macroeconomics: dynamical systems and dynamic optimization, and familiarity with RBC and New Keynesian models.

Part I builds the theory of dynamical systems and solution methods from the ground up, so students who have already studied difference equations and log-linearization will find the early classes familiar. No prior exposure to Bayesian methods, to sequence-space methods, or to heterogeneous-agent models is assumed. Programming is in the Dynare modelling language, in Helix's YAML specification, and in Julia or MATLAB for the labs. Prior programming experience is helpful but not required.

Course structure

Part	Focus	Sessions
Opening	What an economic model is, and what this course is for	Session 1
Part I — Solution methods	The mathematics of DSGE models: dynamical systems and solution methods, on the whiteboard	Sessions 2–8
Part II — Model building	Dynamic programming and model building: RBC and New Keynesian models, by hand and in Dynare and Helix	Sessions 9–14
Part III — Estimation	Taking models to data: Bayesian and limited-information methods	Sessions 15–20
Part IV — Heterogeneity	Incomplete markets and HANK models in the sequence space	Sessions 21–27
Closing	End-term project presentations and wrap-up	Session 28

Assessment

The grade is based on problem sets, two projects, and a final exam:

Component	Weight
Problem sets (four, one per part; 5% each)	20%
Mid-term project — estimating a DSGE model	25%

Component	Weight
End-term project — a heterogeneous-agent model in Helix	30%
Final exam	25%

Problem sets are one per part; the Part I set is analytical, and the rest are primarily computational, building directly on the lab work. The final exam is written and covers the entire course. The two projects remain the spine of the course:

- **Mid-term project (assigned Session 18).** Extend and estimate a published small-scale DSGE model in Dynare on real data, reproduce its key results.
- **End-term project (assigned Session 24).** Build, document and validate a heterogeneous-agent model in Helix — either a replication of a published result or an original extension. Complete YAML specification with a short accompanying write-up. Projects of sufficient quality may be published in the Helix model library under the student's name, with full credit; see the note on Helix below.

Session schedule

Note: Provisional and may be adjusted as the semester progresses.

Part I — The mathematics of DSGE models (Sessions 2–8)

Session	Topic	Reading
2	Dynamical systems I — deterministic difference equations: scalar first- and second-order linear equations, lag operators, first-order linear systems (nonsingular and singular), stability, phase diagrams, and a first look at nonlinear systems.	Miao, ch. 1
3	Dynamical systems II — stochastic difference equations: information sets and conditional expectations, first-order linear stochastic systems, scalar linear rational-expectations models, the method of undetermined coefficients and lag-operator methods.	Miao, ch. 2 (2.1–2.2)
4	Solving multivariate linear rational-expectations systems: the Blanchard–Kahn method and conditions — determinacy, indeterminacy and sunspots; Klein's method and Sims' gensys; the generalized Schur (QZ) decomposition; the recursive law of motion; eigenvalues and saddle-path stability.	Miao, ch. 2 (2.3–2.4)

Session	Topic	Reading
5	From economic models to dynamical systems: the canonical form $E_t\, f(y_{t+1}, y_t, y_{t-1}, \varepsilon_t) = 0$; the deterministic steady state and how to find it; linearization — Taylor expansion and log-linearization.	Fernández-Villaverde et al. (2016)
6	Reading a solved model: impulse responses, the MA(∞) representation, unconditional moments via the Lyapunov equation, variance and historical shock decompositions, simulation and filtering.	Fernández-Villaverde et al. (2016)
7	Perturbation beyond first order: the policy function as an unknown object, second- and third-order expansions, risk and precautionary terms, the shift in the ergodic mean, pruning and GIRFs. And, when local methods fail, a brief survey: value-function and policy iteration, projection and collocation, occasionally binding constraints and the ZLB — deliberately a tour, with references for further study.	Fernández-Villaverde et al. (2016)
8	Sequence-space methods (representative agent): the model as H(U, Z) = 0 on a truncated horizon; Newton on stacked time paths; the sequence-space Jacobian; the IRF as a numerical derivative; equivalence with the state-space solution at first order.	Auclert et al. (2021)

Part II — Dynamic programming and model building (Sessions 9–14)

Session	Topic	Hands-on / lab
9	Dynamic programming, in brief: sequential versus recursive formulations, the Bellman equation and the principle of optimality, value and policy functions, Euler equations and transversality conditions — just enough to build the models that follow. (Miao, chs. 6–7)	Solve a finite-horizon consumption–saving problem by backward induction, by hand and in code.
10	The RBC model, by hand: environment and primitives, first-order conditions, recursive competitive equilibrium, the steady state, log-linearization — and solving the linearized system with the tools of Part I.	Derive the RBC first-order conditions and steady state by hand; solve the steady state numerically (Newton, homotopy); check the Blanchard–Kahn conditions.

Session	Topic	Hands-on / lab
11	From economics to code: the Dynare model language — timing conventions, model and steady-state blocks, shocks and stoch_simul; diagnosing failures (Blanchard–Kahn violations, singularity, steady-state errors). Helix: the same model in YAML.	Build the RBC model from scratch in Dynare and in Helix; verify that the two solutions agree — the state-space / sequence-space equivalence, made concrete.
12	The New Keynesian model I: nominal rigidity (Calvo and Rotemberg), the New Keynesian Phillips curve, the dynamic IS curve, Taylor rules, the Taylor principle and determinacy — derived by hand, then taken to the computer.	Build the three-equation New Keynesian model; map the determinacy region over the policy-rule coefficients.
13	The New Keynesian model II: a medium-scale variant — capital and investment adjustment costs, habit formation, price and wage indexation, sticky wages, a richer shock structure. Calibration.	Extend the baseline to medium scale; compare impulse responses across model variants.
14	Transitory versus permanent, and beyond: shocks that return to the same steady state versus transitions between steady states; MIT shocks; perfect foresight as stacked-time Newton — the direct payoff to Session 8; long-run growth, demographic and fiscal transitions. Model-building synthesis.	A permanent fiscal reform solved with Dynare's perfect-foresight solver and in Helix; open clinic — students write, solve and debug a model of their own choosing, with support.

Part III — Estimation (Sessions 15–20)

Session	Topic	Hands-on / lab
15	From solution to likelihood: state-space form, observation equations, the choice and treatment of observables, detrending, stochastic singularity and measurement error.	Write the observation equations for the New Keynesian model; assemble and inspect a real dataset.
16	The Kalman filter and the likelihood: prediction and update recursions, the prediction-error decomposition, initialization; a note on particle filtering for non-linear solutions.	Code the Kalman filter from scratch; evaluate and plot the likelihood in one and two parameters.
17	Bayesian inference and identification: priors and prior predictive checks, the posterior, mode-finding and the Laplace approximation; local identification, weak identification and observational equivalence.	Dynare's identification toolbox; find the posterior mode; diagnose a deliberately non-identified parameter.

Session	Topic	Hands-on / lab
18	Posterior simulation: MCMC, random-walk Metropolis–Hastings, proposal scaling and acceptance rates, burn-in and multiple chains, convergence diagnostics; the marginal likelihood and model comparison. Brief note on sequential Monte Carlo.	Launch a full estimation — completed as a take-home given the runtime. Mid-term project assigned.
19	Using the estimates: posterior IRFs with credible bands, smoothed shocks, historical and forecast-error variance decompositions, conditional forecasts and counterfactuals; practical pitfalls.	Read the output of the take-home estimation; produce decompositions and a forecast.
20	Limited-information estimation: GMM and SMM moment matching, impulse-response matching to identified VAR or local-projection responses, minimum distance and the weighting matrix, inference, and robustness to misspecification.	IRF-match the New Keynesian model to estimated VAR responses.

Part IV — Heterogeneous-agent models (Sessions 21–27)

Session	Topic	Hands-on / lab
21	Why heterogeneity: incomplete markets and idiosyncratic risk, the failure of aggregation, the evidence on marginal propensities to consume, and the Bewley–Huggett–Aiyagari class of models.	MPCs and wealth distributions in the data, set against the representative-agent benchmark.
22	The household problem I: discretizing income processes (Tauchen, Rouwenhorst, and why Rouwenhorst for persistent processes), asset grids and interpolation; value function iteration, policy iteration and the endogenous gridpoint method; borrowing constraints.	Implement the endogenous gridpoint method for a consumption–savings problem.
23	The household problem II: the stationary distribution — the transition operator, histogram and lottery methods, eigenvector versus iteration; Aiyagari general equilibrium and the capital supply curve.	A full stationary Aiyagari equilibrium in Helix.

Session	Topic	Hands-on / lab
24	Aggregate dynamics: Krusell–Smith and approximate aggregation versus the sequence-space approach; the heterogeneous-agent Jacobian — expectation vectors and the fake-news algorithm; assembling the general-equilibrium Jacobian from blocks.	Compute and visualize household Jacobians in Helix; read their columns as impulse responses.
25	One-asset HANK: monetary transmission with heterogeneous MPCs, direct versus indirect effects, and the intertemporal Keynesian cross.	Build a one-asset HANK model in Helix; decompose a monetary policy shock.
26	Two-asset HANK: liquid and illiquid assets, the wealthy hand-to-mouth, portfolio adjustment costs; fiscal policy, transfers and multipliers; the distributional incidence of inflation and fiscal policy.	A two-asset HANK model; a transfer experiment and its distributional consequences.
27	Frontier and synthesis: estimating heterogeneous-agent models (the sequence-space likelihood, IRF matching in HANK), non-linearity and occasionally binding constraints, and open questions in the field.	End-term project clinic.

Session 28 — Closing

End-term project presentations and course wrap-up.

Software and computing

Parts II and III use Dynare, which runs on MATLAB or on the free GNU Octave. Octave is a fully workable free alternative for everything the course requires. Parts II and IV use Helix, an open-source solver for DSGE and heterogeneous-agent models that works in the sequence space. Helix takes model input as YAML and requires no programming, though the labs will show that a small amount of scripting makes life easier. It is available in the browser and as a Julia package.

A note on Helix

Helix is under active development, and students in this course will be among its first users. As a result, not every part of the tool will be finished in time. Bayesian estimation in Helix in particular may not be ready during this term. Part III is therefore taught entirely in Dynare, which is mature and well documented. The course does not depend on Helix for anything it cannot already do reliably.

Additionally, students in this course will help shape the tool. Bug reports, feature requests, and honest assessments are welcome. The end-term project is built with this in mind: a well-documented, validated Helix model may be published in the public Helix model library under its author's name. For students, this is a citable, public artefact of their own work — useful in PhD applications and job applications alike. For the project, it expands the library of worked examples that make the tool useful to everyone else.

Recommended bibliography

Part I — Solution methods

Miao, J. (2020). Economic Dynamics in Discrete Time, 2nd ed. MIT Press. (chapters 1–2 are the primary reference for Sessions 2–4; chapters 6–7 for the dynamic programming of Session 9)

Fernández-Villaverde, J., J. Rubio-Ramírez and F. Schorfheide (2016). "Solution and Estimation Methods for DSGE Models." Handbook of Macroeconomics, Vol. 2.

Auclert, A., B. Bardóczy, M. Rognlie and L. Straub (2021). "Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models." Econometrica 89(5), 2375–2408.

Part II and III — Model building and Estimation

Galí, J. (2015). Monetary Policy, Inflation, and the Business Cycle, 2nd ed. Princeton University Press. (primary text for the New Keynesian sessions)

Christiano, L., M. Eichenbaum and C. Evans (2005). "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." Journal of Political Economy 113(1), 1–45.

Christiano, L., M. Trabandt and K. Walentin (2010). "DSGE Models for Monetary Policy Analysis." Handbook of Monetary Economics, Vol. 3, 285–367.

Dynare Reference Manual, and Helix documentation and worked examples (course platform).

Part IV — Heterogeneous-agent macroeconomics

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

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

Kaplan, G., B. Moll and G. Violante (2018). "Monetary Policy According to HANK." American Economic Review 108(3), 697–743.

Helix documentation, worked examples, and the public model library (course platform).