Firm Dynamics and Productivity Growth

Session 2: Canonical Models of Firm Dynamics

Pedro Armada, Ph.D.

Halle Institute for Economic Research
Friedrich Schiller University Jena

June 17, 2026

Contents

Why study firm dynamics

Common Ancestry

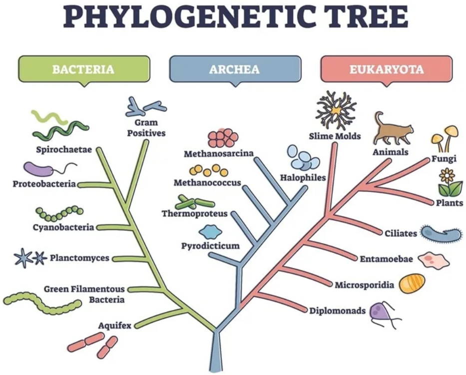

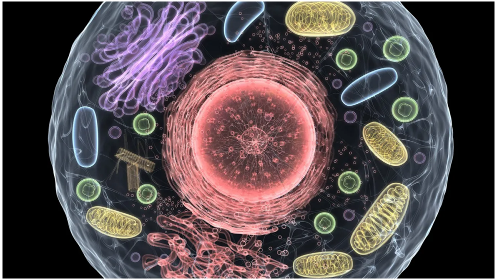

Just as biology traces complex organisms back to common ancestors, macroeconomics traces many models of production back to the **representative firm**.

Why study firm dynamics

- Much of macroeconomics starts from a **representative firm**.

- This is useful because it keeps the aggregate problem simple:

$$Y_t = F(K_t, L_t)$$

- But it hides the fact that production is carried out by many different firms.

- Firm dynamics asks what we lose when we compress all firms into one aggregate producer.

Why the fiction works

The representative-firm setup works because, under strong assumptions, different ownership structures give the same aggregate allocation:

- households rent capital to firms;
- households directly own firms;
- households own firms through stocks.

The same Euler equation and aggregate resource constraint emerge.

See notes: Mukoyama, "Capital Accumulation by Firms".

Where the fiction breaks

The equivalence breaks once firms are not all the same or face different constraints.

- Some firms are more productive than others.
- Some firms are young, others are mature.
- Some firms grow, others exit.
- Some firms innovate, others do not.
- Some firms can borrow, others are financially constrained.
- Some firms invest smoothly, others adjust rarely and in large spikes.

Beyond the representative firm

Firm dynamics moves from

$$Y_t = F(K_t, L_t)$$

to

$$Y_t = \int y_t(i)\, di$$

Aggregate outcomes depend on the **distribution of firms**, not just on aggregate capital and labor.

- Aggregate productivity depends on which firms grow, shrink, enter, and exit.
- Aggregate dynamics depend on how shocks are distributed across firms.

Where to start

Where to start

- Firm dynamics sits at the intersection of macroeconomics, trade, IO, labor, finance, development, etc.

- Modern firm dynamics is built around a tight link between **micro data** and **quantitative theory**.
    - empirical facts motivate the assumptions, and
    - mechanisms in the model explain the moments.

- **Heterogeneity** is the new standard.

Heterogeneous agent models

AEA Papers and Proceedings 2025, 115: 168–171
https://doi.org/10.1257/pandp.20251117

Heterogeneous Agent Models[†]

By Ayşe İmrohoroğlu[*]

Heterogeneous agent models now lie at the core of most research in macroeconomics. In many applications, they have replaced the representative agent framework. However, both use what is in the "super" core of macroeconomics, which is the use of microfoundations. Most modern macroeconomic models use frameworks composed of optimizing agents, firms, and a government. Some use partial equilibrium models, others general equilibrium. There is not one model that fits all questions. Depending on the question, we generate models with different components. Almost all of them, however, start with optimization by some agents in the economy.

Why firm heterogeneity matters

- Firms differ substantially even within narrowly defined industries:
  - sales, employment, assets;
  - productivity, innovation, investment;
  - access to credit, international markets.
- These differences are persistent and correlated across dimensions, suggesting that firm heterogeneity reflects underlying economic mechanisms.

Prototype of a quantitative project

Question $\rightarrow$ Theory $\rightarrow$ Computation $\rightarrow$ Insight

1. **Model specification:** define agents, preferences, technology, frictions, and equilibrium.
2. **Numerical solution:** compute agents' decisions and equilibrium objects.
3. **Calibration / validation:** discipline parameters using data and check model fit.
4. **Model analysis:** run counterfactuals, policy experiments, and welfare analysis.

1. Model Specification

- Specify agents: households, firms, government, external sector.
- Define preferences and technologies:

$$u(c_t, \ell_t), \qquad y_t = A_t f(k_t, n_t).$$

- Introduce frictions: borrowing limits, adjustment costs, fixed costs, shocks.
- State the equilibrium concept: partial/general equilibrium, competitive/Nash equilibrium, etc.

2. Numerical Solution

Once the model is specified, solve agents' optimization problems and impose equilibrium conditions.

Static problems:

- Examples: input choices, pricing decisions.
- Typical methods: first-order conditions, root-finding, constrained optimization.

Dynamic problems:

- Examples: investment, entry and exit.
- Typical methods: value function iteration, policy function iteration, Euler equation methods.

3. Calibration / Validation

Calibration: choose parameters so that simulated moments match empirical targets.

- Target moments should correspond to the mechanisms the model is designed to explain.

Validation: compare model outcomes with data beyond the calibration targets.

- Useful checks include untargeted moments and distributions.

4. Model Analysis

If the model explains well what we can observe, we can use it to study what we cannot observe.

- Conduct policy experiments and counterfactuals: taxes, subsidies, shocks, reforms.
- Compare steady states or transition paths.
- Study distributional effects across firms, workers, sectors, or cohorts.
- Evaluate welfare and aggregate implications.

Where to find data

- **Kauffman Firm Survey**: U.S.; startups; 2004–2011; panel.
- **Survey of Consumer Finances**: U.S.; hhs/bus. owners; 1989–2022; rep. cross-section.
- **Compustat** (commercial): U.S.; listed firms; panel.
- **Orbis** (commercial): several countries; private and public firms; panel.
- **Annual Survey of Industries**: India; manufacturing; 1975–2024; rep. cross-section.
- **World Bank Enterprise Surveys**: 160+ countries; rep. cross-sections.
- **Banco de Portugal Central Balance Sheet**: Portugal; firms; 2006–2023; panel.
- **Banco de España Microdata on Individual Enterprises**: Spain; firms; 1995–2024; panel.
- **CompNet**: Europe; micro-aggregated firm distributions.

Technology and big data in economics

AEA Papers and Proceedings 2020, 110: 42–48
https://doi.org/10.1257/pandp.20201058

EMPIRICAL PRACTICE IN ECONOMICS: CHALLENGES AND OPPORTUNITIES‡

Technology and Big Data Are Changing Economics: Mining Text to Track Methods†

By Janet Currie, Henrik Kleven, and Esmée Zwiers*

Technology and big data in economics

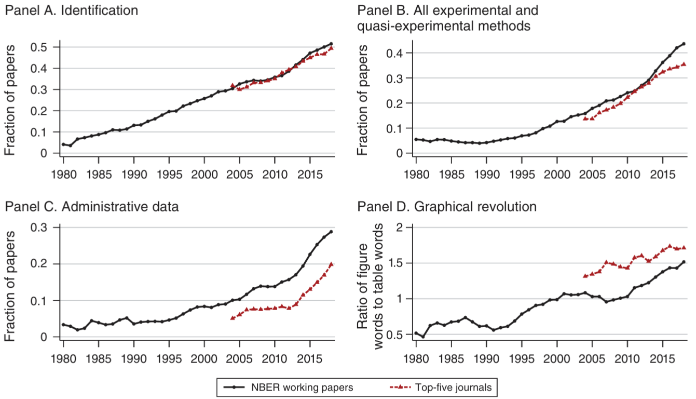

Figure 2. The Credibility Revolution

Notes: This figure shows different dimensions of the "credibility revolution" in economics: identification (panel A), all experimental and quasi-experimental methods (panel B), administrative data (panel C), and the graphical revolution (panel D). Panel D shows the ratio of the number of "figure" terms to the number of "table" terms mentioned. See Table A.I for a list of terms. The series show five-year moving averages.

Where to find code

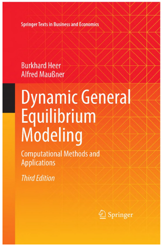

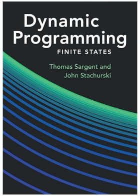

Macroeconomic Model Data Base

QuantEcon

Econ ARK

OPENICPSR

HARVARD Dataverse

Choosing a programming language

Do I still need to learn how to code in the age of AI?

- **Yes!** LLMs can be helpful assistants, but we must still be the architects and supervisors.
- To use AI intelligently, we must understand how the code works.

Why Python?

- General-purpose, open-source language.
- Free, widely supported, and driven by an open community.
- Combines simplicity with powerful scientific libraries.

Python dominance

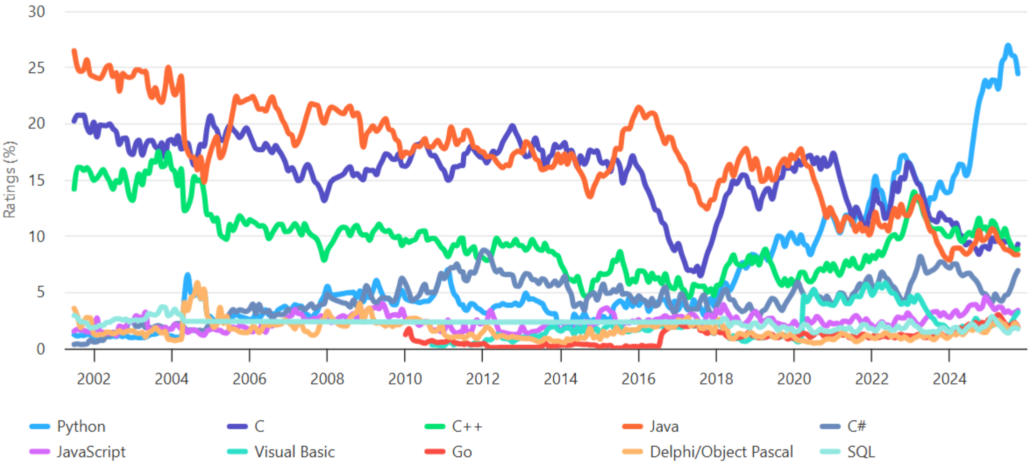

Figure: TIOBE Programming Community Index

Python dominance

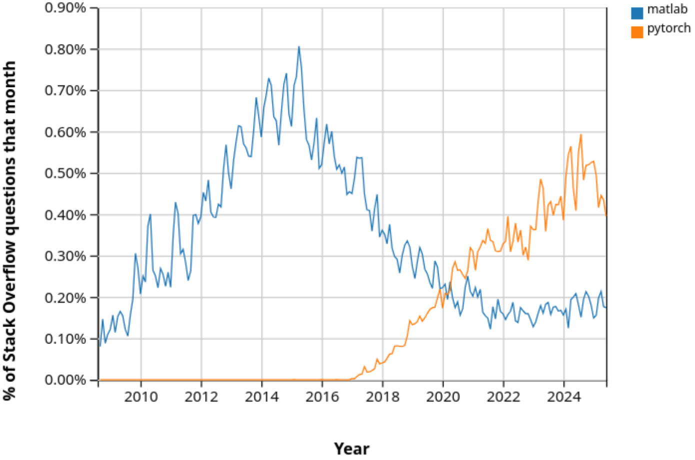

Figure: QuantEcon

Comparing programming languages

Matlab, Python, Julia: What to Choose in Economics?

Chase Coleman[1] · Spencer Lyon[1] · Lilia Maliar[2] · Serguei Maliar[3]

Accepted: 8 April 2020

Abstract
We perform a comparison of Matlab, Python and Julia as programming languages to be used for implementing global nonlinear solution techniques. We consider two popular applications: a neoclassical growth model and a new Keynesian model. The goal of our analysis is twofold: First, it is aimed at helping researchers in economics choose the programming language that is best suited to their applications and, if needed, help them transit from one programming language to another. Second, our collections of routines can be viewed as a toolbox with a special emphasis on techniques for dealing with high dimensional economic problems. We provide the routines in the three languages for constructing random and quasi-random grids, low-cost monomial integration, various global solution methods, routines for checking the accuracy of the solutions as well as examples of parallelization. Our global solution methods are not only accurate but also fast. Solving a new Keynesian model with eight state variables only takes a few seconds, even in the presence of an active zero lower bound on nominal interest rates. This speed is important because it allows the model to be solved repeatedly as would be required for estimation.

Keywords Toolkit · Dynamic model · New Keynesian model · Global nonlinear · Low discrepancy · Quasi Monte Carlo

Comparing programming languages

Matlab, Python, Julia: What to Choose in Economics?

Table 1 Accuracy and speed of value iterative methods

Degree	L_1	L_∞	Julia CPU (s)	Matlab CPU (s)	Python CPU (s)
Conventional VFI (VFI)					
2nd	− 3.83	− 2.76	1.05	30.94	10.97
3rd	− 4.97	− 3.32	0.92	20.85	6.67
4th	− 6.06	− 4.03	1.13	18.02	5.80
5th	− 7.00	− 4.70	1.00	13.80	4.51
Envelope condition method (ECM)					
2nd	− 3.83	− 2.76	0.29	0.32	0.37
3rd	− 4.97	− 3.32	0.21	0.21	0.24
4th	− 6.06	− 4.03	0.27	0.22	0.27
5th	− 7.00	− 4.70	0.14	0.10	0.15
Endogenous grid method (EGM)					
2nd	− 3.81	− 2.76	0.35	27.19	5.05
3rd	− 4.95	− 3.34	0.25	18.51	3.43
4th	− 6.06	− 4.05	0.29	13.95	3.19
5th	− 7.04	− 4.73	0.19	10.43	2.28

[a] L_1 and L_∞ are, respectively, the average and maximum of absolute residuals across optimality conditions and test points (in log 10) units on a stochastic simulation of 10,000 observations. CPU is the time necessary for computing a solution (in s)

Canonical models of firm dynamics

Canonical models of firm dynamics

- **Productivity and selection:** Hopenhayn (1992)

  Firms differ in productivity. Entry and exit generate a stationary distribution of firms.

- **Age and capital accumulation:** Clementi–Palazzo (2016)

  Firms accumulate capital gradually, so age matters even conditional on size.

- **Financial frictions:** Cooley–Quadrini (2001); Clementi–Hopenhayn (2006); Albuquerque–Hopenhayn (2004)

  Firms may be unable to finance efficient scale. Net worth and contracts determine growth and survival.

Canonical models of firm dynamics

- **Investment frictions:** Caballero–Engel (1999); Khan–Thomas (2008)

  Fixed adjustment costs generate inaction, lumpy investment, and nonlinear aggregate responses.

- **Innovation:** Klette–Kortum (2004)

  Firms grow by adding product lines and shrink when competitors displace them.

See notes: Roldan, "Notes on Firm Dynamics".

In this session we will focus on Hopenhayn (1992) and its numerical implementation.

Hopenhayn (1992)

How do entry, exit, and idiosyncratic productivity shocks determine the long-run equilibrium of an industry?

- Competitive industry with a continuum of firms.
- Firms differ in productivity φ.
- Firms face persistent idiosyncratic shocks.
- Low-productivity firms exit, new firms enter.
- The distribution of firms is an equilibrium object.

Aggregates are stationary. No aggregate uncertainty.

Environment

- Firms produce a homogeneous good using one input: labor.
- Output and input markets are competitive.
- Firm productivity is $\varphi \in S = [0, 1]$.
- Production is: $q = f(\varphi, n)$.
- Productivity follows a Markov process: $F(d\varphi' \mid \varphi)$.
- Operating firms pay a fixed cost c_f.
- Entrants pay a sunk entry cost c_e.

Static firm problem

Given prices (p, w), a firm with productivity φ chooses labor:

$$\pi(\varphi, p, w) = \max_{n \geq 0} \{pf(\varphi, n) - wn\} .$$

This gives three static objects:

$$\pi(\varphi, p, w), \qquad q(\varphi, p, w), \qquad n(\varphi, p, w).$$

The fixed operating cost c_f is paid after production.

Why? Without a fixed operating cost, low-productivity firms would not exit.

Timing

For an **incumbent** firm:

1. Productivity φ is observed.
2. The firm chooses labor n.
3. The firm produces and earns $\pi(\varphi, p, w) - c_f$.
4. The firm decides whether to continue or exit.

Entrants pay c_e, draw initial productivity from ν, and then operate as incumbents.

Value of exit is normalized to zero.

Incumbent problem

In a stationary equilibrium with prices (p, w), the incumbent value is

$$v(\varphi) = \pi(\varphi, p, w) - c_f + \beta \max \left\{ 0, \int v(\varphi')F(d\varphi'|\varphi) \right\}.$$

The firm **continues** if

$$\int v(\varphi')F(d\varphi'|\varphi) \geq 0.$$

Exit is characterized by a cutoff x:

$$\begin{cases} \text{exit} & \text{if } \varphi < x, \\ \text{continue} & \text{if } \varphi \geq x. \end{cases}$$

Entry

Potential **entrants** are ex ante identical. After paying the sunk entry cost c_e, entrants draw initial productivity from ν.

The expected value of entry is

$$v^e = \int v(\varphi)\nu(d\varphi).$$

Free entry implies

$$v^e \leq c_e,$$

with equality whenever entry is positive:

$$M > 0 \quad \Rightarrow \quad v^e = c_e.$$

Industry distribution

Let μ denote the measure of firms over productivity.

For any set $A \subseteq S$, $\mu(A)$ is the mass of firms with productivity in A.

Aggregate output supply is

$$Q^s(\mu, p, w) = \int q(\varphi, p, w)\mu(d\varphi).$$

Aggregate labor demand is

$$N^d(\mu, p, w) = \int n(\varphi, p, w)\mu(d\varphi).$$

Prices clear the output and labor markets:

$$p = D(Q), \qquad w = W(N).$$

Law of motion

Let M denote the mass of entrants.

For any measurable set $A \subseteq S$:

$$\mu'(A) = \underbrace{\int_{\varphi \geq x} F(A|\varphi)\mu(d\varphi)}_{\text{surviving incumbents}} + \underbrace{M\nu(A)}_{\text{new entrants}} .$$

In a stationary equilibrium: $\mu' = \mu = \mu^*$.

Stationary equilibrium

A stationary equilibrium consists of

$$p^*, w^*, Q^*, N^*, x^*, M^*, \mu^*$$

such that:

1. Firms choose labor optimally.
2. Incumbents choose exit optimally.
3. Free entry holds.
4. Output and labor markets clear.
5. The firm distribution is stationary:

$$\mu^*(A) = \int_{\varphi \geq x^*} F(A|\varphi)\mu^*(d\varphi) + M^*\nu(A).$$

Mechanism

The model generates firm dynamics through **selection**.

- Firms receive idiosyncratic productivity shocks.
- More productive firms are larger and more likely to survive.
- Low-productivity firms exit.
- Entrants replace exiting firms.

Aggregate productivity depends on which firms survive.

Discrete approximation

To solve the model, we replace the continuous state space with a productivity grid:

$$\varphi_1, \ldots, \varphi_n.$$

For a given price p, the Bellman equation is approximated by:

$$v_i = \pi(\varphi_i, p) + \beta \max \left\{ 0, \mathbb{E}\left[v(\varphi' | \varphi_i) \right] \right\}.$$

In practice, expectations are computed using simulation or numerical integration.

Computing the stationary distribution

Given the exit cutoff $\bar{\varphi}$, firms evolve as:

$$\varphi_{t+1} = \begin{cases} \varphi', & \text{if } \varphi_t \geq \bar{\varphi}, \\ \varphi^{\mathsf{e}}, & \text{if } \varphi_t < \bar{\varphi}. \end{cases}$$

Exiting firms are replaced by new entrants drawn from the entrant distribution.

Simulating many firms over time gives an approximation to the stationary distribution.

Live demo

QuantEcon

Quantitative Economics with Python using JAX

Thomas J. Sargent and John Stachurski

Last changed: Jun 06, 2024 ▾

19. The Hopenhayn Entry-Exit Model

GPU

This lecture was built using a machine with access to a GPU — although it will also run without one.

Google Colab has a free tier with GPUs that you can access as follows:

1. Click on the "play" icon top right
2. Select Colab
3. Set the runtime environment to include a GPU

19.1. Outline

The Hopenhayn (1992, ECMA) entry-exit model is a highly influential (partial equilibrium) heterogeneous agent model where

- the agents receiving idiosyncratic shocks are firms, and

On this page

Powered by Jupyter Book

Where to go from here

Where to go from here

- **Financial frictions**

  Alfaro–Bloom–Lin (2024); Guo (2024); Li–Lian–Ma–Martell (2025); Boar–Gorea–Midrigan (2026); Peter (2026); Ebsim et al. (2026); Tan (2026).

- **Innovation, intangibles, market power**

  Crouzet et al. (2022); Aghion et al. (2023); Crouzet–Eberly (2023); Aghion et al. (2023); De Ridder (2024); He–Mostrom–Sufi (2025); Hampole et al. (2025); Cullen et al. (2025); Yotzov et al. (2026); Baslandze et al. (2026).

- **Markup dispersion**

  Edmond–Midrigan–Xu (2023); Raval (2023); Baqaee–Farhi–Sangani (2024); Döpper et al. (2025); Burstein–Carvalho–Grassi (2025); Alviarez et al. (2025); Benkard et al. (2025).

Where to go from here

- **Startup heterogeneity**

  Sterk–Sedláček–Pugsley (2021); Hopenhayn–Neira–Singhania (2022); Bernard et al. (2022); Karahan–Pugsley–Şahin (2024); Bailey et al. (2026); Dinlersoz et al. (2026).

- **Worker-firm dynamics**

  Bilal et al. (2022); Elsby–Gottfries (2022); Berger–Herkenhoff–Mongey (2022); Queiró (2022); Hurst–Kehoe–Pastorino–Winberry (2025).

- **Production networks**

  Carvalho et al. (2021); Elliott–Golub–Leduc (2022); Alfaro-Ureña–Manelici–Vasquez (2022).