

UC **SANTA BARBARA**
Laboratory for Aggregate
Economics and Finance

Volume XIX, No 1, September 2026

IN THIS ISSUE

St. Louis Federal Reserve and Washington University LAEF Macro-Labor Conference

Conference held April 2–3, 2026

Market Power Conference

Conference held April 10–11, 2026

5th Annual Workshop on Labor Markets and Macroeconomic Outcomes

Conference held April 17–18, 2025

Artificial Intelligence and the Macroeconomy

Conference held April 24–25, 2026

Demographics in the 21st Century

Conference held April 30 - May 1, 2026



FROM THE LAB

1. Assistant Director's Message - *Nick Pretnar*

St. Louis Fed–WashU–LAEF Macro-Labor Conference

Note: speakers are highlighted in author listings

3. Subjective Uncertainty and the Marginal Propensity to Consume
Gizem Koşar and Davide Melcangi

3. Imperfect Information and Slow Recoveries in the Labor Market
Anushka Mitra

4. Job Transformation, Specialization, and the Labor Market Effects of AI
Lukas B. Freund and Lukas Mann

4. Generative AI in the Labor Market: Evidence and Some Theory
Adam Blandin

5. The Labor Market Return to Permanent Residency
Kory Kroft, Isaac Norwich, Matthew J. Notowidigdo, and Stephen Tino

5. The Labor Demand and Labor Supply Channels of Monetary Policy
Sebastian Graves, Christopher Huckfeldt, and Eric T. Swanson

6. Job Search, Wages, and Inflation
Laura Pilossoph and Jane M. Ryngaert

6. New Evidence on Hiring Costs
Fabian Lange and David Wiczer

Market Power

Note: speakers are highlighted in author listings

8. Trends in Average Market Power: Where Do We Stand?
Chad Syverson

10. Technology-Driven Market Concentration through Idea Allocation
Yueyuan Ma and Shaoshuang Yang

12. Ideas and Firm Dynamics When It Takes Two to Tango
Seula Kim, Jane Olmstead-Rumsey, and Honghao Wang

14. Scaling Up: How Technology and Policy Shape Firm Dynamics
Thomas Hasenzagl

- 16. Input Price Dispersion Across Buyers: New Evidence and Implications for Aggregate Productivity
Ariel Burstein, Javier Cravino, and Marco Rojas
- 18. The Rise of Specialized Financial Products
Ana Babus, Matias Marzani, and Sara Moreira
- 20. Natural Centralization in Decentralized Finance
Pablo Azar, Adrian Casillas, and Maryam Farboodi
- 22. Consumer Credit Regulation and Lender Market Power
Zachary Bethune, Joaquin Saldain, and Eric R. Young
- 24. The Local Origins of Business Formation: Entry as a Two-Stage Process
Emin Dinlersoz, Timothy Dunne, John Haltiwanger, and Veronika Penciakova
- 26. Monopsony Power and the Transmission of Monetary Policy
Bence Bardóczy, Gideon Bornstein, and Sergio Salgado
- 28. Robinson Meets Roy: Monopsony Power and Comparative Advantage Mark
Bils, Barış Kaymak, and Kai-Jie Wu
- 30. Product Market Power as a Labor Distortion: the Case of Home Production Substitutes
William Jungerman, Stan Rabinovich, and Bailey Sanders
- 32. Pass-through with Price Dispersion
Brian C. Albrecht and Mark Whitmeyer

5th Annual Workshop on Labor Markets and Macroeconomic Outcomes

Note: speakers are highlighted in author listings

- 35. Spousal Insurance Around the World
Annika Bacher, Kevin Donovan, Philipp Grübener, and Todd Schoellman
- 35. The Job Ladder, Unemployment Risk, and Incomplete Markets
Alex Clymo, Piotr Denderski, Yusuf Mercan, and Benjamin Schoefer
- 36. Signals, Interviews, and Hiring in Labor Market Equilibrium
Piero Gottardi, Benjamin Lester, Ryan Michaels, and Ronald Wolthoff
- 36. Reassessing Labor Market Flows: The Role of Marginal Workers and the Composition of Quits and Layoffs
Kathrin Ellieroth and Amanda Michaud
- 37. Taxing the Rich? A Theory of Income and Wealth Inequality
V.V. Chari, Patrick Kehoe, Pierlauro Lopez, and Elena Pastorino

37. Home Bias and the Gender Wage Gap

Christian Hellwig

38. Job Transformation, Specialization, and the Labor Market Effects of AI

Lukas B. Freund and Lukas Mann

38. Masculinity Norms, Occupational Choice, and Reallocation: Micro and Macro Evidence

Rachel Schuh and Martina Uccioli

39. Self-employment as Self-insurance

Daniel Jaar

39. The Changing Value of Employment

Giovanni Gallipoli and Davide Alonzo

40. Hiring with a Quantity-Skill Tradeoff

Fil Babalievsky, Henry Hyatt, and Lawrence Warren

Artificial Intelligence and the Macroeconomy

Note: speakers are highlighted in author listings

42. Simulating the Survey of Professional Forecasters

Anne Lundgaard Hansen, John J. Horton, Sophia Kazinnik, Daniela Puzzello, and Ali Zarifhonarvar

42. Did I Make Myself Clear? The Fed and the Market under the 2020 Monetary Policy Framework

Anna Cieslak, Michael McMahon, and Hao Pang

43. Identifying Episodes of Fiscal Austerity: An LLM-Based Approach

Karan Bhasin and Prakash Loungani

43. Seniority Bias, Vintage Mismatch, and Productivity in the AI Era

Aakash Kalyani

44. Measuring Task-Level Technological Exposure: A Language Model Approach

Andre Mouton

44. AI in the Office and the Factory: Evidence from Administrative Software Registry Data

Gustavo de Souza

45. AI-Augmented Capital-Skill Complementarity

Andre Victor Luduvic

45. Understanding Growth Through Automation: The Neoclassical Perspective

Lukasz Drozd, Mathieu Taschereau-Dumouchel, and Marina Mendes Tavares

46. The Macroeconomic Implications of the Gen-AI Economy

Pablo Guerrón-Quintana, Tomoaki Mikami, and Jaromir Nosal

46. Monetary Policy Shocks: A New Hope — Large Language Models and Central Bank Communication

Ruben Fernandez Fuertes

47. Deep Learning in the Sequence Space

Marlon Azinovic-Yang and Jan Žemlička

47. A Global Solution Method for HACT Models with Aggregate Risk

Niklas Bonnmann

Demographics in the 21st Century

Note: speakers are highlighted in author listings

49. Understanding Latin America's Fertility Decline: Age, Education, and Cohort Dynamics

Pablo Guerrón-Quintana, Tomoaki Mikami, and Jaromir Nosal

49. The Effect of Cognitive Skills on Fertility Timing

Agustin Diaz Casanueva

50. Financial Incentives to Fertility: From Short to Long Run

Lidia Cruces and F. Javier Rodríguez-Román

50. Childcare and Long-Term Care: Altruism and Strategic Motives within the Family

Chao Fu, Ami Ko, and Chamna Yoon

51. Boomerang College Kids: Unemployment, Job Mismatch and Coresidence

Stefania Albanesi, Rania Gihleb, and Ning Zhang

51. The Demographic Cliff and the Market for Higher Education: Implications for Public Finance

Irina Popova

52. Parental Leave Policies, Fertility, and Labor Supply

Daisoon Kim and Minchul Yum

52. The Power of Policy Incentives: Female Labor Supply, Fertility and Parental Leave Policy Design

Mary Ann Bronson

53. Marriage and Work Among Prime-Age Men

Adam Blandin, John Bailey Jones, and Fang Yang

53. Optimal Pension Reforms during Demographic Transitions

Selahattin Imrohoroglu

54. Population Aging, Firms, and Worker Composition

Santiago Caicedo, Cheng Chen, Aspen Gorry, and Takahiro Hattori

54. The Elasticity of Social Security Policy

Kathleen McKiernan

55. Redistribution, Distortions, and the Welfare Effects of Social Security

Youngsoo Jang, Svetlana Pashchenko, and Ponpoje Porapakkarm

55. Why Aging Induces Deflation and Secular Stagnation

R. Anton Braun and Daisuke Ikeda

56. The Postpandemic US Immigration Surge: New Facts and Inflationary Implications

*Anton Cheremukhin, Sewon Hur, Ronald Mau, Karel Mertens, Alexander W. Richter,
and Xiaoqing Zhou*

From the Lab
**Message from the
Assistant Director**
Nick Pretnar



This is the first edition of the nineteenth volume of From the Lab. As we start the new 2026-2027 academic year, we look back on five consecutive conferences we featured in spring of 2026. These conferences all took place in five different corners of the United States, with Santa Barbara anchoring the mid-point of LAEF's busy spring schedule with the April Market Power Conference. We also returned to the St. Louis Fed and Wash U in St. Louis for what has become a fun annual tradition. Similarly, our relationship with Vanderbilt continues to blossom, as we held the 5th Annual Workshop on Labor Markets and Macroeconomic Outcomes in Nashville in April 2026, and we plan to return this conference to Santa Barbara in April 2027 as part of a renewed, three-year commitment with Vanderbilt University. Look out for the call for papers for the 6th annual version to come later this fall.

This issue of From the Lab features 50 paper abstracts and summaries from the 50 talks and seminars that LAEF helped highlight in its spring conferences. This means that from our efforts and through the generous funding we receive from UCSB, LAEF was able to give 50 researchers from all corners of quantitative macroeconomic theory venues to showcase and workshop 50 of their papers. In these seminars they received feedback from myriad attendees. Thus, LAEF helped push the frontier of research forward along 50 different dimensions, and perhaps many more, in just a few months this spring. This just highlights how broad our impact is with regards to the conference services we provide.

Partnering with other institutions to put on conferences (as we do now more than ever) is the greatest service we provide to the profession, and one that, in my travels, I am often told is very much appreciated. At LAEF we make efforts to always try to have open calls, and to reserve some slots for young researchers who may or may not have access to a rather rigid and hierarchical professional seminar structure, but for whom feedback's marginal benefits are enormous. We look forward to continuing our mission to advance the frontier of knowledge for this 2026-2027 academic year, too.



St. Louis Federal Reserve and Washington University LAEF Macro-Labor Conference

April 2–3, 2026; Federal Reserve Bank of St. Louis and Washington University in St. Louis, MO

Attendees:

Davide Melcangi — *Federal Reserve Bank of New York*
Anushka Mitra — *Federal Reserve Board*
Lukas B. Freund — *Boston College*
Adam Blandin — *Vanderbilt University*
Matthew J. Notowidigdo — *University of Chicago*
Christopher Huckfeldt — *Federal Reserve Board*
Jane M. Ryngaert — *University of Notre Dame*
Fabian Lange — *McGill University*

Organizers:

Alexander Bick — *Federal Reserve Bank of St. Louis*
Serdar Birinci — *Federal Reserve Bank of St. Louis*
Victoria Gregory — *Federal Reserve Bank of St. Louis*
Serdar Ozkan — *Federal Reserve Bank of St. Louis*
Philipp Grübener — *Washington University in St. Louis*



Subjective Uncertainty and the Marginal Propensity to Consume

Presenter: *Davide Melcangi*

Coauthors: *Gizem Koşar*

Earnings uncertainty is central to most heterogeneous-household models, yet there is surprisingly little evidence on how subjective uncertainty relates to consumption behavior. Using data from the Survey of Consumer Expectations, the authors show that the marginal propensity to consume (MPC) is increasing and concave in individual-specific earnings growth uncertainty — the opposite of what a standard consumption-savings model with risk heterogeneity implies. The paper pinpoints which model mechanisms create this disconnect and shows how recently proposed deviations from the full-information rational expectations framework can reconcile theory with the empirical findings.



Imperfect Information and Slow Recoveries in the Labor Market

Presenter: *Anushka Mitra*

The unemployment rate remains elevated long after recessions, a persistence that standard search-and-matching models cannot explain. Noise shocks — expectational errors due to noise in received signals about aggregate shocks — account for much of this sluggishness: using a structural VAR, the author finds that absent noise shocks, unemployment would have recovered to its pre-recession level six quarters earlier over 1968–2019. A search-and-matching model with on-the-job search, endogenous search effort, and wage rigidity shows imperfect information generates two channels of persistence — slow learning amplifies persistent productivity shocks, and noise shocks provide an additional source of sluggishness, magnified by sticky wages and vacancy posting.



Job Transformation, Specialization, and the Labor Market Effects of AI

Presenter: *Lukas B. Freund*

Coauthors: *Lukas Mann*

Who will gain and who will lose as AI automates tasks? While much discourse focuses on job displacement, the paper shows job transformation — a shift in the task content of jobs — creates large and heterogeneous earnings effects. The authors develop a quantitative, task-based model in which occupations bundle multiple tasks and workers with heterogeneous task-specific skills select into occupations by comparative advantage; automation shifts the relative importance of tasks within each occupation, inducing wage effects characterized analytically. Measuring task content with natural language processing and projecting automation effects from large language models, the authors find moderate exposure benefits workers on average but high exposure harms them, with large within-occupation dispersion; the return to social skills rises while the return to analytical skills falls; and low-earners gain more than high-earners.



Generative AI in the Labor Market: Evidence and Some Theory

Presenter: *Adam Blandin*

How Between 1 and 5 percent of all work hours are currently assisted by generative AI, with workers reporting time savings equivalent to about 1.4 percent of total work hours — evidence that substantial productivity gains from generative AI are possible. At the macro level, industries with higher AI adoption rates have experienced faster productivity growth (though causality is not established), a relationship of similar magnitude in Europe and the U.S.; there is no clear evidence that industry-level AI adoption is associated with employment changes.



The Labor Market Return to Permanent Residency

Presenter: *Matthew J. Notowidigdo*

Coauthors: *Kory Kroft, Isaac Norwich, and Stephen Tino*

Many temporary foreign worker programs issue “closed” visas that tie workers to a single employer, restricting mobility and bargaining power. Using administrative data linking matched employer-employee records in Canada to temporary and permanent visa records from 2004–2014, the authors study the labor market return to gaining permanent residency (PR), which loosens this restriction. An event-study design shows gaining PR leads to a sharp, immediate, and persistent increase in the job-switching rate of 21.7 percentage points and a 5.7 percent increase in earnings three years after PR; workers also sort into higher-wage firms, with the increase in the firm pay premium accounting for roughly 56 percent of the total earnings gain.



The Labor Demand and Labor Supply Channels of Monetary Policy

Presenter: *Christopher Huckfeldt*

Coauthors: *Sebastian Graves, and Eric T. Swanson*

Monetary policy is conventionally understood to influence labor demand with little effect on labor supply. Estimating the response of labor market flows to high-frequency changes in interest rates around FOMC announcements and Fed Chair speeches, the authors find that — contrary to the consensus view — a contractionary monetary policy shock leads to a significant increase in labor supply: workers reduce the rate at which they quit jobs to nonemployment, and non-employed individuals increase job-seeking behavior. These effects are quantitatively important: holding supply-driven labor market flows constant, the decline in employment from a contractionary monetary policy shock would be twice as large. A heterogeneous agent model with frictional labor markets and an active labor supply margin is estimated to interpret the findings.



Job Search, Wages, and Inflation

Presenter: *Jane M. Ryngaert*

Coauthors: *Laura Pilossoph*

A canonical job search model incorporating nominal wage rigidities implies that on-the-job search should increase with expected inflation. In a novel survey, the authors show workers are indeed more likely to search at higher values of hypothetical inflation. A back-of-the-envelope calculation suggests the search mechanism and the 2021–2023 rise in inflation expectations can explain roughly 30 percent of the change in the job-to-job transition probability.



New Evidence on Hiring Costs

Presenter: *Fabian Lange*

Coauthors: *David Wiczer*



Market Power

April 10–11, 2026; The Upham Hotel, Santa Barbara, CA

Attendees:

Chad Syverson — <i>University of Chicago, Booth School of Business</i>	Joaquin Saldain — <i>Bank of Canada</i>
Shaoshuang Yang — <i>Chinese University of Hong Kong</i>	John Haltiwanger — <i>University of Maryland</i>
Shenzhen Seula Kim — <i>Penn State University</i>	Sergio Salgado — <i>The Wharton School, University of Pennsylvania</i>
Thomas Hasenzagl — <i>University of Minnesota</i>	Baris Kaymak — <i>Université de Montréal</i>
Ariel Burstein — <i>University of California, Los Angeles</i>	William Jungerman — <i>University of North Carolina at Chapel Hill</i>
Ana Babus — <i>Washington University in St. Louis</i>	Mark Whitmeyer — <i>Arizona State University</i>
Pablo Azar — <i>Federal Reserve Bank of New York</i>	

Organizers:

Nick Pretnar — *LAEF, UCSB*
 Brian Albrecht — *International Center for Law & Economics*
 Thomas Phelan — *Federal Reserve Bank of Cleveland*
 Yueyuan Ma — *University of California, Santa Barbara*



Trends in Average Market Power: Where Do We Stand?

Presenter: *Chad Syverson*

Economists have long suspected that American business has grown less competitive, and a cluster of broad trends seems to point that way. Profit rates have risen, investment has fallen, business dynamism has declined, productivity growth has slowed, the share of income going to labor has shrunk, and common ownership has spread across rival firms. It is not unreasonable, as Chad Syverson noted, to look at this collage of facts and infer that market power has risen across the economy. Syverson presented a survey of the recent literature on exactly that question, tied to an article he wrote the previous year for the *Annual Review of Economics*. He stressed at the outset that there is no settled consensus and that he would lay out competing views rather than dictate a conclusion, and he deliberately declined to draw antitrust prescriptions on the grounds that competition policy is about specific markets, not aggregates.

The paper that galvanized the field was by De Loecker, Eeckhout, and Unger, which used what is called the production approach to measuring markups, a technique tracing to Robert Hall's work in the late 1980s. This method exploits cost minimization: for any input flexible enough to be adjusted to its optimal level each period, the markup is the input's output elasticity relative to its revenue share. Syverson offered an intuition, the elasticity captures what the input lets the firm sell while the revenue share captures what the firm pays for it, so a gap between the two is the markup. Its appeal is that it needs only ordinary production data. Applied to public-firm data, the method produced a striking result, namely that average markups of roughly twenty percent above cost around 1980 had tripled to about sixty percent by the mid-2010s.

Most of the talk worked through the substantial body of critiques that has since accumulated, nearly all of which push the increase downward. Because data usually contain revenue rather than physical output, Syverson explained, an estimate built on a revenue elasticity is mechanically pushed toward a value implying no markup at all, so the method may not measure what it claims. The choice of which input counts as flexible matters, as James Traina showed, and broader definitions yield far smaller increases. Mert Demirer pointed out that the ratio computed for energy versus materials ought to agree under cost minimization but rarely does, suggesting that markups and productivity cannot be cleanly separated from this kind of data. Chris Conlon and others showed that aggregation choices matter, and Syverson called the original revenue-weighting a genuine

mistake, since a true aggregate markup must be cost-weighted. He treated these not as nitpicks but as carrying collective weight, and pointed to two findings from outside the framework: Conlon finds no correlation between an industry's measured markup growth and its price inflation, and work with Ryan Decker finds the link between markups and business dynamism running the wrong way relative to the motivating story.

A deeper conceptual point is that the markup ratio assumes away employer power. Once a firm has monopsony power, the correct expression also involves the elasticity of labor supply to that firm, so 1 the wedge attributed entirely to product markets may partly, or even mostly, reflect markdowns on workers. An audience member pressed that if monopsony comes from search-and-matching frictions, labor is not a flexible static input at all, undermining the whole approach, a point Syverson fully conceded. Another participant noted that market definition matters far more for labor than for tradable products, and Syverson observed that concentration at the narrow local level has if anything declined slightly even as mega-firms have grown. On common ownership, where passive giants hold rival firms, theory due to Rotemberg suggests muted competition, but Syverson stressed there is no smoking gun, no memo telling firms to compete less, only indirect channels such as weakened performance-based executive pay.

His own reading is cautiously revisionist. Markups probably rose, he guessed, but well under a tripling, perhaps as much on the employer side as the product side. And even a rising wedge need not indict antitrust: if technology raised fixed costs and returns to scale, markups may simply have to rise in equilibrium so firms can cover those costs, a technology story rather than a policy failure.



Technology-Driven Market Concentration through Idea Allocation

Presenter: *Shaoshuang Yang*

Coauthors: *Yueyuan Ma*

This paper studies whether changes in technological opportunities can explain two major developments in the U.S. economy since the early 2000s: declining business dynamism and rising market concentration. Existing explanations typically emphasize strategic behavior by incumbent firms, including increased market power, barriers to entry, or anti-competitive conduct. The paper proposes a complementary explanation. Rather than focusing on firms' strategic actions, it argues that both trends can emerge as equilibrium responses to changes in the underlying technological environment.

The central idea is that technological breakthroughs occur in waves. Periods characterized by major innovations create opportunities for new firms to enter and commercialize novel ideas. As the intensity of technological breakthroughs declines, innovative activity becomes increasingly concentrated within incumbent firms, reducing entry and increasing concentration. The paper therefore asks whether fluctuations in technological opportunity can jointly account for the evolution of growth, firm dynamics, and market structure.

The paper begins by documenting two well-known trends in the U.S. economy: aggregate and productivity growth have slowed over time, and market concentration has increased substantially across many industries. At the microeconomic level, these aggregate patterns coincide with a decline in firm entry and an increase in the size and dominance of highly productive incumbent firms. To study the role of technological opportunities, the authors construct measures of "technology waves" using textual measures of patent novelty, patent importance indices, and citation-based measures of technological significance. Breakthrough activity exhibits clear cyclical patterns, with major peaks corresponding to periods such as the biotechnology expansion of the 1980s and the information technology boom of the late 1990s, though the overall intensity of breakthroughs has declined over time. The authors show that these breakthrough measures are negatively related to future market concentration, motivating the idea-allocation mechanism at the core of the paper.

The theoretical framework focuses on inventors who must decide whether to commercialize their ideas through startups or incumbent firms. The model introduces two key features of incumbent firms: they possess complementary assets, accumulated knowledge, and organizational capabilities that generate synergies when commercializing innovations, but they also face adoption frictions, since highly novel technologies are often difficult to integrate into existing production structures. This

trade-off generates a natural allocation mechanism. During periods of intense technological breakthroughs, adoption frictions become relatively more costly, so high-quality ideas are more likely to be commercialized through startups, increasing entry and reducing concentration. When breakthrough intensity declines, incumbents become relatively more effective at commercializing innovation, so entry falls, concentration rises, and aggregate growth slows. The model also generates positive sorting between idea quality and firm characteristics, with larger firms attracting certain types of innovation while startups play a particularly important role during periods of major technological disruption.

Empirically, the authors combine patent data, firm-level information, and measures of market concentration, addressing the concern that breakthroughs may be endogenous to economic conditions with a shift-share instrument that exploits technological diffusion across technology classes. The results support the model's central predictions: technological breakthrough measures are robustly and negatively associated with market concentration and positively associated with innovative activity among young firms, and high-quality innovations are increasingly observed within larger incumbent firms as breakthrough intensity declines. Quantitatively, simulations suggest that declining breakthrough intensity can explain a substantial share of the observed rise in concentration and decline in business dynamism, generating lower entry rates, slower aggregate growth, and greater concentration alongside higher average firm quality.

The paper contributes to a growing literature linking innovation, firm dynamics, and market structure. Its primary contribution is to introduce an idea-allocation channel through which changes in technological opportunities affect concentration. Rather than viewing concentration solely as the result of strategic behavior by incumbents, the paper emphasizes how shifts in the nature of innovation itself can reshape the distribution of firms and ideas. By connecting technological waves, startup formation, concentration, and growth within a unified framework, the paper provides a coherent explanation for the changing structure of the U.S. economy.



Ideas and Firm Dynamics When It Takes Two to Tango

Presenter: *Seula Kim*

Coauthors: *Jane Olmstead-Rumsey, and Honghao Wang*

This paper studies how the organization of innovative activity within firms affects firm dynamics and market structure. The motivation comes from a dramatic rise in collaborative invention over the last several decades. In the United States, team patents have become increasingly dominant, rising from roughly forty percent of all patents in the 1970s to nearly eighty percent in recent years. At the same time, innovative activity has become increasingly concentrated within large firms.

The paper asks whether these two trends are connected. More specifically, it investigates how team formation among inventors influences the allocation of innovative labor, firm growth, and concentration. While a large literature studies firm heterogeneity and innovation, relatively little attention has been paid to the role of team formation as a mechanism shaping firm dynamics.

The paper begins by documenting the rise of team-based innovation using U.S. patent data. Most patents are now generated by teams rather than individual inventors, and collaborative invention has become the dominant form of innovative production, even as inventive activity has become increasingly concentrated within large firms. An important observation is that most inventor collaborations occur within firms rather than across organizational boundaries, so firms serve not only as production units but also as environments in which inventors find collaborators and form teams. The empirical evidence further shows that inventive concentration has increased even more rapidly than employment concentration, suggesting that changes in the organization of innovation may be an important factor behind broader changes in market structure.

The central innovation of the paper is a model of team formation embedded within a firm-dynamics framework. Workers can either innovate individually or collaborate with another inventor, and the productivity of a team depends on the quality of the match between collaborators; firms assign inventors to the collaborations that generate the highest expected productivity. A key implication is that larger firms provide larger pools of potential collaborators, so the expected productivity of team formation rises with firm size, giving larger firms an organizational advantage in generating productive collaborations. This mechanism creates increasing returns to scale in innovative activity, as firms that grow larger attract more inventors, which in turn improves the quality of potential matches and raises innovative productivity, reinforcing concentration in both innovative activity and firm size. The model is embedded within an otherwise standard framework featuring heterogeneous firms, endogenous entry and exit, and knowledge spillovers.

The empirical analysis combines patent records, Census-linked firm information, and inventor-employer matched data covering 1976 to 2018, and the authors estimate an innovation production function relating patent output to inventor employment in order to recover parameters governing returns to scale and the productivity advantages of team production, distinguishing solo invention productivity from team-based productivity over time. The paper finds that productivity associated with solo invention has declined, consistent with the view that ideas are becoming harder to discover, and that average team productivity has also declined — but the relative advantage of team production has increased, implying that collaboration has become increasingly important for generating successful innovation. These changes help explain both the rise of team patents and the growing concentration of inventive activity: as the benefits of collaboration increase, larger firms gain an advantage because they provide access to larger networks of potential collaborators, so inventors become increasingly concentrated in large organizations. Quantitatively, the model successfully reproduces the rise of team-based innovation, increasing inventor concentration, larger average firm size, and higher output, while also predicting a decline in knowledge spillovers as innovative activity concentrates within a smaller set of firms.

The paper contributes to the literature by introducing team formation as a central mechanism linking innovation and firm dynamics. Whereas traditional models typically focus on firm-level productivity differences, this framework emphasizes the role of organizational structure and collaboration opportunities, providing a microfoundation for why larger firms may enjoy increasing returns in innovative activity and a useful bridge between the literatures on innovation, labor allocation, and firm dynamics. More broadly, the paper suggests that changes in how innovation is produced may be as important as changes in the quantity of innovation itself, and that the increasing importance of collaborative invention appears capable of reshaping the distribution of firms, inventors, and economic activity throughout the economy.



Scaling Up: How Technology and Policy Shape Firm Dynamics

Presenter: *Thomas Hasenzagl*

When firms in the same industry combine, regulators and economists alike confront a familiar tension: a merger can hand the combined firm greater market power, letting it raise markups and restrict output, but it can also make the firm more efficient, lowering costs and prices. Thomas Hasenzagl presented work in progress that builds a quantitative general-equilibrium model of mergers to weigh these two forces against each other, with a sharp focus on one specific and underappreciated source of efficiency. The gains he isolates come not from production but from overhead — the selling, general, and administrative functions like human resources, accounting, information technology, and compliance. When two firms merge, each arrives with its own back-office staff, and the combined entity does not need to keep all of them. Hasenzagl invokes the blocked Kroger-Albertsons grocery merger as a familiar case where such administrative consolidation might have occurred.

The mechanism that does the work is increasing returns to scale in overhead: larger firms use less overhead per unit of output, so a merger that doubles output need not double the administrative payroll. Hasenzagl estimates that doubling output requires only roughly sixty-four percent more overhead labor. Crucially, these savings show up in marginal cost, not merely as a fixed cost, so the model generates a genuinely U-shaped cost curve. Costs fall as a small firm grows, until a separate force — managerial diseconomies, in the tradition of Lucas’s span-of-control model, where a manager’s limited attention is stretched thinner as the firm expands — takes over and costs begin to rise. Markups, meanwhile, are allowed to vary with firm size, so a larger combined firm faces less elastic demand and charges more. The model is dynamic, with agents each period choosing whether to run a firm or supply labor, a margin that stands in for entry and exit. Periodically a manager is matched with another firm, and a merger goes through only if it generates gains both parties can share.

A merged firm, in Hasenzagl’s setup, operates the sum of the two firms’ product scopes — a demand shifter capturing breadth, brand, or quality — under the acquirer’s productivity. This summing assumption drew the most sustained pushback. Audience members noted that it is a strong, merger-specific synergy assumption, hard to discipline because the empirical literature offers no consensus on what a typical merger efficiency actually is. One suggested estimating synergies in line with classic merger models and comparing the two approaches. Hasenzagl conceded the assumption carries considerable empirical slush, and agreed that pinning down the quantitative magnitude of merger efficiencies is the central challenge, since the precise location of any optimal policy threshold depends on it. Earlier, another participant, drawing on close work with retail firms and on Chad Syverson’s research on the J-curve, pressed on the large up-front intangible investments

real mergers require to integrate operations; Hasenzagi acknowledged that his model abstracts from these adjustment costs entirely.

To study policy, Hasenzagi then introduces a regulator who can block mergers using a single 1 observable quantity, the ratio of overhead to revenue. Allowing every merger to proceed lowers welfare by about half a percent relative to a world with no mergers, because market-power harm dominates on net. But a rule that permits mergers with high overhead savings while blocking large-on-large deals among already-big firms, where savings are slim and the markup harm large, can raise welfare by up to roughly three and a half percent. Hasenzagi was careful to clarify, under questioning, that this is optimal only within the narrow class of single-threshold rules, not a global first best; a true planner would merge firms until each operated at minimum average cost. He floated an overhead-to-revenue cutoff near twenty percent as a workable rule of thumb.

The most provocative result comes from feeding in the historical rise in returns to scale, drawn from companion work with coauthor Louis Barres. As overhead economies have grown over recent decades, plausibly driven by information technology and intangibles, the model produces both a merger wave and higher markups. Yet the welfare-maximizing antitrust stance loosens rather than tightens, because efficiency gains have grown even faster than the costs of concentration — implying that a rising tide of mergers need not call for a heavier regulatory hand.



Input Price Dispersion Across Buyers: New Evidence and Implications for Aggregate Productivity

Presenter: *Ariel Burstein*

Coauthors: *Javier Cravino, and Marco Rojas*

Professor Javier Cravino of the University of Michigan presented work-in-progress findings on the dispersion of prices that firms pay for the same intermediate inputs, using a novel dataset of electronic invoices issued by Chilean firms. The central motivation was to fill a gap in the existing literature: while there is extensive evidence on how consumer prices vary across retail stores, comparatively little is known about how input prices vary across different purchasing firms buying from the same supplier. Cravino's research asks whether these price differences reflect traditional markup variation, differences in the cost of supplying distinct buyers, or something else — and what the implications are for aggregate efficiency in supply chains.

The empirical foundation is a comprehensive set of electronic invoices collected by the Chilean IRS in 2017, covering all firm-to-firm transactions in the country. A key feature of the data is that many invoices contain detailed internal product codes and descriptions, allowing the researchers to identify products at a level of granularity far exceeding what is available in trade data, where the standard unit of comparison is a tariff line. Cravino restricts the analysis to manufacturing firms whose product codes meet stringent consistency requirements — at least four characters in length, with a one-to-one mapping between code and description — and encodes the identity and location of the selling establishment directly into the product definition. This ensures that price comparisons are made across buyers purchasing the genuinely identical product from the same supplier.

The headline empirical fact is that price dispersion across buyers is both large and pervasive. Across all manufacturing sectors, the gap between the 90th and 10th percentile buyer for the same product averages around 60 log points, with a median standard deviation of approximately 6 log points. This dispersion is stable over time despite significant inflation variation during the sample period. Notably, dispersion is substantially lower across establishments owned by the same firm and in regulated industries such as refined petroleum, suggesting the patterns are not an artifact of measurement noise.

Cravino then examined what predicts these price gaps. Proximity between buyer and seller reduces prices only marginally, and payment terms are not associated with higher prices once buyer size is controlled for. Observable buyer characteristics such as size and transaction volume together explain only 4–8% of price gap variation, while buyer fixed effects alone account for roughly a quarter. Cravino interpreted this as evidence that price differences reflect buyer-specific markups rather than differences in the cost of supplying distinct buyers.

An audience member pushed on this point, asking whether the results should be interpreted as a welfare concern or simply as the outcome of bilateral bargaining that redistributes surplus without creating allocative inefficiency. Cravino acknowledged this as a genuinely important interpretive question. He noted that the absence of evidence for fixed price-quantity menus — buyers frequently purchase different quantities at the same price — makes a pure price-discrimination story less likely, and that the patterns are consistent with a bargaining framework in which some buyers systematically negotiate better deals than others.

A second exchange touched on whether the large unexplained residual simply reflects differences in buyer negotiating ability. Cravino conceded that unobserved heterogeneity in bargaining power is a plausible interpretation, and that no single observable characteristic comes close to accounting for the bulk of the variation.

The quantitative section extended a standard wedge-decomposition framework to incorporate buyer-specific markups alongside product-level markup dispersion. The central finding is that equalizing markups across buyers of the same product increases estimated welfare gains by 30 to 60% relative to equalizing markups only across products — a significant contribution that the prior literature, lacking this type of data, had been unable to capture.



The Rise of Specialized Financial Products

Presenter: *Ana Babus*

Coauthors: *Matias Marzani, and Sara Moreira*

Professor Ana Babus of Washington University in St. Louis presented research examining how the variety of financial products that firms use to raise external funds has evolved over time, and whether the type of products firms issue explains how much capital they are able to raise. The central observation motivating the paper is that between 1985 and 2014, the number of distinct financial products available to firms grew from roughly 150 to nearly 800. Beyond stocks and bonds — the most familiar instruments — firms across the economy have increasingly drawn on a wide and heterogeneous array of securities, and Babus’s research asks what that proliferation means for capital raising.

The empirical foundation is the Securities Data Company’s issues database, which covers the universe of public and private security issuances by U.S. firms from 1970 onward. Restricting attention to non-financial sectors over three decades, the dataset encompasses roughly 18,000 issuers across 850 four-digit SIC sectors, approximately 72,000 issuances, and around \$11 trillion in total proceeds. Babus noted that the average issuer completes only about four issuances over the entire thirty-year period, reflecting how infrequently most firms access capital markets. The analysis aggregates issuances into five-year periods to address this sparsity.

A central empirical focus is distinguishing new financial products — those first observed after 1985 — from older ones. New products account for over 80% of the product types used during the sample period and are adopted across roughly 70% of sectors, yet they represent less than 20% of issuances by count while generating about 30% of total proceeds. Babus argued this disproportion suggests something meaningful about their role: when firms do use new products, they raise funds comparable to what median old products generate per issuance, despite the novelty of the instrument.

To quantify the degree of differentiation across products, Babus constructed a novelty index using text analysis of Investopedia articles describing each financial product. The logic is that financial products function as mappings from issuer risk characteristics into investor cash flows, and the textual descriptions of those contingencies capture how different two products actually are. A cosine similarity measure between any product pair yields a novelty index ranging from zero — essentially identical to an existing product — to one — highly distinct. Babus used this measure both to demonstrate that the 600-plus new products represent meaningful variety rather than relabeled versions of the same instrument, and as a proxy for financial specialization in subsequent analysis.

An extended discussion arose around the novelty index itself. One audience member pushed back on conflating novelty with quality, drawing an analogy to consumer goods where new products can be measured along objective characteristic dimensions. Babus agreed the framing required care, clarifying that novelty here is intended to capture differentiation rather than improvement, and that the index's primary purpose is to establish that meaningful variety exists in the data rather than to rank products by desirability.

A more pointed exchange concerned the model's structure. Babus presented a parsimonious framework in which firms choose among available financial products based on a productivity index — essentially a Sharpe ratio capturing expected payoff per unit of risk — subject to strategic competition with other firms issuing the same product. An audience member challenged whether a mean-variance framework could plausibly generate 750 distinct products, arguing that two securities suffice to span all payoff profiles under standard assumptions. Babus acknowledged the artificial element in the model's investor segmentation assumption — investors in one product type cannot trade others — but defended it as a tractable way to introduce the downward-sloping demand that product differentiation requires, noting that institutional investor specialization provides real-world support for the idea.

The quantitative results show that changes in average product productivity explain roughly two-thirds of the variation in sector-level proceeds, and that this productivity component correlates with both the adoption of newer and more specialized financial products — consistent with the paper's central claim that financial innovation is a meaningful determinant of firms' ability to raise external capital.



Natural Centralization in Decentralized Finance

Presenter: *Pablo Azar,*

Coauthors: *Adrian Casillas, and Maryam Farboodi*

Pablo Azar of the Federal Reserve Bank of New York presented research using decentralized finance — specifically the Ethereum blockchain — as a laboratory for studying how intermediaries, market power, and rents arise endogenously even in markets with nominally free entry. While the paper is grounded in crypto markets, Azar was clear that the broader contribution is to the theory of financial intermediation: the Ethereum block-building market provides unusually clean data to examine how profit shares are determined across layers of a disintermediation chain and what drives concentration within each layer.

The institutional backdrop matters. Azar noted that the intersection of crypto and traditional finance has grown substantially, with major banks and payment platforms — including JPMorgan, Bank of America, and Stripe — developing stablecoin and blockchain products following recent regulatory shifts. He flagged several systemic concerns associated with this growth, including cryptoization risk, the fragility of certain stablecoins that operate like shadow banking vehicles, and evidence from academic work that a portion of intermediation profits in crypto markets can be traced to illicit sources. These concerns motivate why understanding the structure of these markets is consequential beyond the crypto context itself.

The specific market Azar studies is Ethereum's block-building ecosystem. For a transaction to be recorded on Ethereum, it must be included in a block proposed every twelve seconds. Arbitrage opportunities — arising from price discrepancies across decentralized exchanges or from liquidation events in lending protocols — can be worth hundreds of thousands of dollars, but broadcasting such a transaction publicly exposes it to front-running. This creates demand for privacy, and with it, a new class of intermediary called block builders, who bundle private transactions and submit sealed bids to block proposers such as Coinbase or Kraken. The proposer selects the highest bid without seeing the underlying transactions, a separation that also serves to partition legal liability. Azar documented that roughly 77% of blocks involve distinct builders and proposers rather than vertical integration.

The central empirical question is how the builder's profit share varies with the value of private versus public information in a given block. Public transactions are those broadcast to the network before appearing on chain; private transactions are those that appear on chain without prior broadcast. Using data from the Yale Mempool Project, which operates a network node to record transaction propagation, the paper can precisely classify each transaction and construct block-level measures of private and public information value alongside the revenue split between builders and proposers.

The key identification challenge is simultaneity: builders choose which transactions to include, how much to bid, and the composition of public and private transactions all at once. To address this, Azar instruments using hacks of Ethereum protocols and broader financial crises such as the SVB collapse. The first stage reveals a clean separation: hacks generate a surge in private information with no corresponding increase in public information, while crises increase both. An audience member asked for an example of an exclusion restriction violation; Azar offered the case where a builder is also the hacker and exploits their position to influence the proposer directly, thereby affecting the bargaining outcome through channels other than information composition.

The second-stage results show that a higher share of private information significantly increases the builder's profit share, while greater public information substantially reduces it — with the public information effect larger in magnitude than the private information effect. The accompanying bargaining model formalizes these results: builders with private information can afford to be patient in choosing which proposer to transact with, and patience translates directly into bargaining power. Public information eliminates that patience and therefore erodes market power.

Azar concluded that private information raises intermediary market power while public information competes it away — a result with implications well beyond decentralized finance for understanding rent extraction and concentration in financial intermediation more broadly.



Consumer Credit Regulation and Lender Market Power

Presenter: *Javier Cravino*

Coauthors: *Ariel Burstein*

This presentation asks when interest rate caps can help borrowers in consumer credit markets. In a standard competitive model of unsecured borrowing and default, the answer is usually no. If lenders price loans competitively, a high interest rate simply reflects high default risk. A cap then removes some contracts that borrowers may want to take, so it can only reduce their choice set. The paper starts from this benchmark and then changes one important assumption: lenders may have market power because some borrowers do not search well.

The model is a dynamic model of unsecured credit. Households receive income shocks, can save or borrow, and can default. Lenders understand that default risk matters when they price loans. The new ingredient is a partially directed search friction. Some households are informed and can search for the best credit terms. Other households are uninformed and randomly draw one offer from the market. Lenders do not observe a borrower's information type, but they do observe income, wealth, and current financial status.

This information friction creates two types of lending markets in equilibrium. Low-cost lenders post terms to attract informed borrowers, who can compare offers. High-cost lenders post much more extractive terms and hope to meet uninformed borrowers who search randomly. The lenders are not different by assumption; they are ex ante the same, but the search friction gives them different incentives, which is why the model can generate two prices for borrowers with the same default risk. The presentation described this as a pecking-order story: households use cheaper credit first and may end up in the high-cost market as they become more constrained or get rejected elsewhere. High-cost lenders enter more heavily in markets for low-income and low-wealth households, since these households need credit more urgently and have worse outside options, giving high-cost lenders more surplus to extract — an endogenous result, not an assumption imposed by the authors.

The model is calibrated to match both ordinary unsecured credit (negative net worth shares, bankruptcy rates, debt-to-income, credit-card rates) and payday lending (the high average annualized payday rate, the share of households using payday loans, and the persistence of payday borrowing) — with only about 4 percent of households being uninformed. That small fraction is enough to generate a large interest rate gap: informed borrowers in the low-cost market pay about 10.8 percent on average versus 26.1 percent for uninformed borrowers in that same market, while the high-cost market averages about 303.9 percent, close to the payday target of 339 percent. Default risk alone cannot explain payday-style rates without unrealistic default behavior — market power breaks the

tight link between default risk and high interest rates.

The main policy experiment is an interest rate ceiling. The welfare-maximizing cap is around 26–27 percent APR, though caps as low as 20 percent can still raise aggregate welfare. The average welfare gain is small, about 0.03 percent of consumption, but distributional gains are larger: low-income and low-wealth households can gain up to about 0.32 percent of consumption, and in the calibrated model the optimal cap is Pareto-improving, so no group loses relative to the unregulated steady state. The mechanism runs mainly through market composition — when a cap binds on high-cost lenders, those lenders earn less surplus and some exit, so uninformed households are less likely to meet extractive payday-like offers and more likely to meet low-cost lenders. There are also costs: as the cap tightens, borrowing becomes cheaper in some parts of the market and default rises, and if the ceiling falls below roughly 22 percent APR the credit market can collapse because lending is no longer profitable enough relative to default risk and entry costs.

The main takeaway is that the welfare effect of credit regulation depends on why interest rates are high. If high rates only compensate lenders for default risk in a competitive market, caps hurt borrowers. If high rates also reflect lender market power over uninformed borrowers, caps can help. The paper's contribution is to show this in a dynamic credit model where high-cost payday-like lending appears endogenously, where poorer households face worse market composition, and where a simple interest rate ceiling can get close to the welfare gains from making all borrowers informed.



The Local Origins of Business Formation: Entry as a Two-Stage Process

Presenter: John Haltiwanger

Coauthors: Emin Dinlersoz, Timothy Dunne, and
Veronika Penciakova

Most research on business startups observes firms only after they hire their first employee. John Haltiwanger presented research using newly linked administrative microdata from the Census Bureau to track the universe of U.S. business applications from initial filing through their transition into employer firms. The authors model business entry as a two-stage process. First, startups come up with a business idea (proxied by an EIN application). Afterwards, the firm enters a gestational period which ends with the hiring of its first worker. The latter is the typical definition of startup formation in the literature.

The talk then decomposed spatial variation in startup intensity (number of startups per capita) into its two components: the number of EIN applications which captures the idea generation phase and the transition rate of applications into an employer startup. The authors find that both margins vary substantially across space, are only weakly correlated, and respond differently to local conditions. In general, high-startup locations are characterized by elevated application rates, while startup deserts are distinguished by especially low transition rates.

Within their universe of observed firms, the authors focus on applications where the applicant indicates plans to pay wages. About 60 percent of these never become employer businesses within eight quarters, confirming that applying for an EIN marks the start of a distinct gestational phase rather than a final bureaucratic step. A participant asked whether the application address reliably captures where the business will operate, which is important given the paper's focus on explaining spatial variation in startup intensity. Haltiwanger reassuringly noted that over 90 percent of application addresses match the eventual business address at the county level.

To formalize the notion of a two-stage business entry process, Haltiwanger presented a model where potential entrepreneurs first decide whether to pursue an idea, captured by paying a cost to enter a gestational phase, and then, after receiving informational signals about the viability of the business, decide whether to hire employees. A participant asked what distinguishes the two stages in practice and whether entrepreneurs could gather information before applying. Haltiwanger emphasized that empirically, the data show a multi-quarter nascent phase, and that the two margins respond to local conditions differently. This represents a pattern inconsistent with the canonical one-stage model of business entry. Another participant asked whether policy frictions such as bureaucratic red tape would affect the transition stage. Haltiwanger confirmed that such frictions would shift the distribution of signals during the gestational phase, lowering transition

probabilities in the process.

A variance decomposition shows that roughly two-thirds of tract-level variation in startups per capita reflects variation in applications, while one-third reflects transition rates. Ranking tracts by startup intensity reveals a striking asymmetry: superstar tracts stand out through high application rates, whereas the weakest tracts ("startup deserts") are dragged down by especially low transitions. The case of African American population share illustrates why this decomposition matters. It is well known that areas with higher African American shares have lower startup intensity. The two-stage framework reveals offsetting forces: higher African American share is associated with high applications per capita but substantially lower transition rates. There is no shortage of entrepreneurial ideas, but the challenge appears to lie in converting them into operating businesses. Haltiwanger noted that adverse local financial market conditions are a plausible explanation for this result, which the authors are exploring in ongoing work. A participant asked whether superstar and desert tracts could be identified by observables, with Haltiwanger noting that local covariates indeed account for a substantial portion of the ranking patterns.



Monopsony Power and the Transmission of Monetary Policy

Presenter: *Sergio Salgado*

Coauthors: *Bence Bardóczy, and Gideon Bornstein*

Over the past four decades, local labor markets in the United States have become less concentrated. Yet, fewer than 6 percent of firms still account for 40 percent of private-sector employment. This paper asks whether a firm's local labor market power shapes the transmission of monetary policy in terms of its wage bill and employment. The answer is yes: firms with high local labor market power respond significantly less to monetary shocks. Using administrative Census data and a New Keynesian model with oligopsonistic competition, Sergio Salgado and his coauthors show that firms with low monopsony power respond about 50 percent more to an expansionary shock than their high monopsony power counterparts. Additionally, they demonstrate that incomplete passthrough reduces the aggregate output response by roughly a quarter, and that the long decline in local firm concentration since 1980 has amplified the effectiveness of monetary policy by more than 50 percent.

Their empirical analysis draws on the Census Longitudinal Employer-Household Dynamics (LEHD) dataset merged with the Longitudinal Business Database (LBD). In total, the data covers firms in 23 states from 1993 to 2021. Importantly, the authors classify a firm as having high monopsony power if it accounts for more than 10 percent of the wage bill in a local labor market, defined as a county by 3-digit NAICS subsector cell. Using local projections to identify the impact of monetary policy shocks on a firm's wage bill, Salgado finds that for a 25 basis point reduction in the federal funds rate, the wage bill of firms with low monopsony power responds roughly 50 percent more than that of firms with high monopsony power within the same local market. During the talk, a participant asked whether this differential might simply reflect the fact that smaller firms tend to be more financially sensitive. The presenter noted that controlling for firm age, national employment size, and even product-market concentration (monopoly power) leaves the core finding largely intact, though the magnitude shrinks modestly when size is absorbed.

To interpret these findings, the authors develop a heterogeneous agent New Keynesian model in which a continuum of local labor markets is populated by a finite number of firms competing à la Bertrand with Calvo-style wage rigidities. A key theoretical result is that with fully flexible wages, the presence of oligopsonistic markdowns does not generate heterogeneous pass-through. In this case, all firms adjust wages proportionally, their market shares stay constant, and there is no differential response. Salgado shows that it is the combination of sticky wages and convex markdown functions that is required to qualitatively replicate the empirical findings.

In particular, the model highlights two channels through which oligopsony dampens aggregate output. First, a partial passthrough channel exists, in which firms optimally raise their markdowns when demand expands rather than fully transmitting higher prices to wages. A lower passthrough of wages reduces the overall employment response and consequently, the firm's output response. Second, a misallocation channel occurs, in which the employment share shifts toward less productive low-monopsony firms after a monetary expansion. This has the important effect of pulling down aggregate TFP since these smaller firms are less productive on average. Quantitatively, these two channels together reduce the output response to a monetary shock by about 24-30 percent relative to a constant-markdown benchmark, with partial passthrough accounting for roughly two-thirds of the gap and misallocation for one-third. A participant asked whether the decline in local labor market concentration at the national level is being confounded with other trends at the aggregate level. The presenter clarified that the historical counterfactual holds aggregate structure fixed and recalibrates only the local firm-size distribution.



Robinson Meets Roy: Monopsony Power and Comparative Advantage

Presenter: *Barış Kaymak*

Coauthors: *Mark Bils, and Kai-Jie Wu*

This presentation gives a different way to think about monopsony power in labor markets. A lot of the recent discussion connects monopsony to large employers, local labor market concentration, and low-wage workers. The paper does not reject that view, but it argues that it misses an important source of employer power: worker-firm match productivity. A worker may be especially productive at one firm relative to other possible firms. That relative productivity gives the worker a good match with the firm, but it also gives the firm room to pay less than the worker's productivity. The main point is that monopsony power can come from comparative advantage, not only from employer size.

The paper combines Robinson's model of non-competitive wage setting with Roy's model of worker sorting. Each worker has a productivity level at each firm; firms observe the worker's productivity at their own firm but not elsewhere, post wage schedules based on productivity, and workers choose the employer with the highest offer, generating an upward-sloping labor supply curve for each firm. The key object is the wage markdown — one minus the ratio of wage to productivity — which here is not assumed but comes from the joint distribution of a worker's productivity across firms: a worker highly productive at one firm but not elsewhere has weak outside options, letting that firm mark the wage down more. So the largest markdowns appear where comparative advantage is strongest, meaning high-productivity workers can face larger markdowns if comparative and absolute advantage align — one of the more surprising parts of the paper.

The theoretical model formalizes this with a symmetric setup of K firms in which each worker's productivity across firms follows a joint distribution described with copulas. Workers go to whichever firm they're best matched to, so the allocation is efficient even though wages sit below productivity, and as the number of competing firms rises, or as productivities become more correlated across firms, outside options improve and markdowns fall. Weaker correlation among high-productivity workers implies larger top-end markdowns; stronger correlation implies smaller ones. An asymmetric extension shows a generally more productive firm pays higher wages, grows larger, and marks down more — the familiar large-firm story — while a merely more skill-biased firm marks down its high-skill workers more heavily without necessarily being large, distorting sorting away from the efficient allocation.

The empirical part uses Brazilian administrative data (RAIS) on formal-sector worker-establishment matches, wages, hours, occupations, industries, and locations, defining labor markets by metropolitan area and occupation and splitting workers within each employer into wage quartiles. Comparing

how wages and employment respond, for workers at different points of an employer's wage distribution, when that employer expands identifies the labor supply elasticity facing the employer and its implied markdown. High-wage workers within a firm face much larger markdowns — wages rise more for them as employers grow, while employment responds more at the bottom — implying markdowns near zero in the bottom wage quartile but around 25 percent or more at the top. High-wage employers, more than large employers, drive this pattern: large employers alone do not face meaningfully less elastic labor supply, weakening the view that size or local concentration is the main source of monopsony power in the Brazilian data.

The final part of the paper studies minimum wages and informality. Because firms interact strategically, a minimum wage shifts the whole wage schedule upward and reduces markdowns even for workers earning above the minimum, and an informal sector gives workers an outside option that can raise formal wages through competition. The policy result is mixed: a higher minimum wage can raise wages and formal employment for more productive workers by reducing markdowns and pulling them out of informal jobs, but it can also push less productive workers into lower-paying informal work.

The main lesson is that monopsony power is not only about large employers paying low wages in concentrated markets — it can also come from some workers being especially good matches for particular firms, which creates rents that firms tax through wage markdowns. In Brazil, that appears to mean high-wage workers within firms and workers at high-wage firms, leading to the paper's somewhat unexpected conclusion: monopsony power may reduce within-firm wage inequality even while it distorts worker sorting and reduces the pay of workers who create the most value for their employers.



Product Market Power as a Labor Distortion: the Case of Home Production Substitutes

Presenter: *William Jungerman*

Coauthors: *Stan Rabinovich, and Bailey Sanders*

Despite decades of convergence, significant gender gaps remain in labor force participation. A large literature has identified falling costs of home production substitutes such as appliances and childcare as a key driver of rising female labor supply, since women historically bear the bulk of home production. However, these studies universally treat the relevant product markets as perfectly competitive, assuming cost reductions fully pass through to prices. William Jungerman and his coauthors argue this assumption is an important omission and introduce monopolistic competition in the market for home production substitutes into a simple model of household time allocation.

In their model, a household consists of a husband, who works full-time in the market, and a wife, who divides her time between market work and home production. Home goods can be produced using a combination of the wife's time and purchased inputs, where each input variety is supplied by a monopolistic firm. A participant asked early on whether specific market services would be assumed to substitute for specific home activities, or whether this would be identified in the data. Jungerman clarified that the model treats purchased inputs as combining with the household's own time to produce each variety of home good, so the price of each input directly affects how much time the household devotes to that activity.

Two key elasticities govern the results of the model: the elasticity of substitution between time and purchased inputs within each home good variety, and the elasticity of substitution across different product varieties. The sign of the difference between these elasticities determines whether endogenous markups dampen or amplify the effects of falling input costs and narrowing gender wage gaps on female labor supply. A participant questioned whether there was an additional channel of imperfect substitutability between husband and wife due to productivity differences in home production. Jungerman acknowledged this but noted that the model assigns home production to one household member, abstracting from within-couple bargaining.

To quantify the model, the authors exploit cross-sectional variation across household types using data from Nielsen, the Consumer Expenditure Survey, the American Time Use Survey, and the Current Population Survey. Their identification strategy backs out parameters sequentially from observed cost shares and expenditure ratios. One participant raised concerns about error propagation from this step-by-step approach and suggested estimating the model jointly on microdata via a MLE estimator. Jungerman noted that they had also run a method-of-moments estimation and were exploring robustness checks, but emphasized the appeal of the clean closed-form mapping between data and parameters. A participant questioned the estimated cross-variety elasticity of around 14 and argued that unobserved heterogeneity could be driving the result.

In quantitative terms, Jungerman's findings suggest that markups on home production substitutes cost households about one percent of consumption, which corresponds to roughly \$650 per year. Furthermore, he estimates that ignoring the endogenous markup response understates the labor supply effects of cost reductions by about 10 percent, while the amplification of wage convergence effects is more modest at around 1 percent. The authors trace this gap to the smaller direct effect of wage changes on home production time relative to cost changes.



Pass-through with Price Dispersion

Presenter: Mark Whitmeyer

Coauthors: Brian C. Albrecht

When economists ask how a tax or an input-cost shock gets divided between sellers and buyers, they usually have in mind a market with a single going price that can be nudged up or down. But many of the markets people actually shop in do not have a price at all. Gas stations along the same road post different numbers, online listings for the same item scatter widely, and price-comparison sites exist precisely because identical goods sell for many different amounts at once. Mark Whitmeyer presented work, joint with Brian Albrecht, that asks how cost shocks pass through to prices in dispersed markets, and his central message was a reassuring one: although limited consideration replaces the single price with a whole distribution of prices, the analysis of pass-through remains entirely tractable. If the audience left with one thing, he said, it should be that price dispersion, at least as he and Albrecht conceive of it, causes no trouble.

The engine of the model is what Whitmeyer called limited consideration. Rather than comparing every seller, each consumer looks only at some subset of firms and buys the cheapest among those she considers, provided the price falls below a reservation value. Because different consumers consider different subsets, the same good ends up selling at many prices. This generates a useful distinction between captive demand, the buyers who consider only one firm and so cannot be lured away, and contested demand, the buyers shared across rivals. In this setting market power is defined behaviorally: a firm has power when it can hold a high margin without losing much business, which amounts to being insulated from direct price comparison. The framework builds on the consideration-set oligopoly of Armstrong and Vickers, adding downward-sloping demand, and the relevant cost shocks are fuel taxes, wholesale changes, and the fees and commissions a platform charges its sellers.

The analytical move at the heart of the talk is to stop thinking about firms as choosing prices and instead think of them as choosing margins. Whitmeyer showed that the pricing game can be rewritten as an equivalent game in normalized margins with the same lowest-price-wins logic, and that in this rewritten game both cost and the curvature of demand simply drop out of the strategic problem. His main theorem is that the equilibrium distribution of margins depends only on the market's consideration structure and not on cost or demand. This separates the problem into a competition layer, where one solves for equilibrium from market structure alone and which is invariant to cost, and a curvature layer, where that cost-invariant margin distribution is mapped back into prices and where all the incidence actually happens.

The clean slogan that curvature does not matter drew the sharpest exchange. An audience member pressed the point that the demand curve a firm actually faces, shaped by the probabilities of the various consideration sets, surely has a curvature that affects its optimal margin. Whitmeyer conceded the slogan was too strong and narrowed it carefully: what washes out is the curvature of an individual consumer's demand, while the shape of residual demand coming from the consideration 1 structure still matters, though it lives entirely within the competition layer and so changes what pass-through is without complicating how it is computed. A second exchange asked how much rides on every firm sharing a common cost and every consumer sharing a common demand curve. Here he drew an asymmetry: demand heterogeneity could be accommodated by building up piecewise, but cost differences remained the genuinely open problem, and he admitted he and Albrecht had not found a non-arbitrary way to handle them.

With the separation in hand, the comparative statics come out intuitively. Highly elastic demand dampens pass-through, since a near-competitive firm earns no rents to pass along, while inelastic demand facilitates it; under perfect competition pass-through is complete. Firms pricing closer to the reservation value, the ones already extracting the most surplus, pass through less. When one audience member argued that the ratio of captive to total reach, rather than reach itself, ought to drive pricing, Whitmeyer agreed the tidy ordering held only under the special case of independent consideration, offering the meta-point that such intuitions partly dissolve because rivals' prices are already responding to the very same forces.



5th Annual Workshop on Labor Markets and Macroeconomic Outcomes

April 17–18, 2025; West End Tower, 17th Floor, Nashville, TN

Attendees:

Michael Waugh — *Federal Reserve Bank of Minneapolis*
 Ryan Michaels — *Federal Reserve Bank of Philadelphia*
 Elena Pastorino — *Stanford University*
 Daniel Jaar — *European University Institute*
 Lukas Mann — *Arizona State University, W.P. Carey School of Business*
 Christian Hellwig — *Toulouse School of Economics*
 Amanda Michaud — *Federal Reserve Bank of Minneapolis*
 Davide Alonzo — *Université de Montréal*
 Todd Schoellman — *Federal Reserve Bank of Minneapolis*
 Rachel Schuh — *Federal Reserve Bank of New York*

Lawrence Warren — *U.S. Census Bureau, Center for Economic Studies*
 Yusuf Mercan — *Federal Reserve Bank of Kansas City*
 Adam Blandin — *Vanderbilt University*
 Kathleen McKiernan — *Vanderbilt University*
 Nick Pretnar — *LAEF, UCSB*
 César Urquizo — *Vanderbilt University*
 Steven Lugauer — *University of Kentucky*
 Giovanni Gallipoli — *University of British Columbia, Vancouver*
 Santiago Lisa — *Carnegie Mellon University*

Organizers:

Adam Blandin — *Vanderbilt University*
 Kathleen McKiernan — *Vanderbilt University*
 Nick Pretnar — *LAEF, UCSB*
 César Urquizo — *Vanderbilt University*



Spousal Insurance Around the World

Presenter: *Todd Schoellman*

Coauthors: *Annika Bacher, Kevin Donovan, Philipp Grübener*

The added worker effect — labor force entry upon spousal job loss — is stronger for young than old households. Using a life-cycle model of two-member households in a frictional labor market that endogenizes asset and human capital accumulation and job offer arrival rates, disciplined against U.S. micro data, the authors study whether this age-dependency is driven by heterogeneous needs for or availability of spousal insurance. Counterfactuals show a strong complementarity: a large added worker effect requires both high spousal earnings potential relative to the primary earner and limited access to other means of self-insurance. The model predicts substantial crowding out of spousal labor supply responses by unemployment benefit extensions among young households, consistent with their stronger insurance motive.



The Job Ladder, Unemployment Risk, and Incomplete Markets

Presenter: *Yusuf Mercan*

Coauthors: *Alex Clymo, Piotr Denderski, and Benjamin Schoefer*

We study a general equilibrium search model with jobs differing in productivity and unemployment risk. Risk-averse workers, searching off and on the job, self-insure against labor income risk via a non-state-contingent bond. On-the-job search mitigates, undoes, or even reverses the risk-aversion-induced inefficiency of search out of unemployment. The interplay of imperfect insurance and unemployment risk heterogeneity creates a new distortion: employed workers climb the job ladder overcautiously. More generous unemployment insurance can increase job creation by kickstarting the job ladder. Quantitatively, these forces improve the equity-efficiency tradeoff and we find the welfare-maximizing replacement rate to be high, in line with current US policy.



Signals, Interviews, and Hiring in Labor Market Equilibrium

Presenter: *Ryan Michaels*

Coauthors: *Piero Gottardi, Benjamin Lester, and Ronald Wolthoff*



Quits, Layoffs, and Labor Supply

Presenter: *Amanda Michaud*

Coauthors: *Kathrin Ellieroth*

We provide fresh evidence that labor supply is broadly counter-cyclical. Employed and laid-off workers both increase labor market attachment when unemployment is high. A workhorse search model is used to frame how these facts add nuance to our understanding of business cycles and of labor supply fundamentals: how many workers are near the margin of participation and what drives their choices? Additional results explore regularities of these patterns across cyclical episodes and in the cross section of workers. Additional microeconomic data that suggest income effects contribute to increased labor supply in recessions but are unlikely to be the complete story.



Taxing the Rich? A Theory of Income and Wealth Inequality

Presenter: *Elena Pastorino*

Coauthors: *V.V. Chari, Patrick Kehoe, Pierlauro Lopez*

Recently, it has been argued that a progressive wealth tax may have a large beneficial impact on the distribution of welfare in society and effectively no adverse effects on real economic activity. This paper evaluates the merits of this view within a dynamic general equilibrium framework that microfound empirically plausible income and wealth distributions as arising from a fundamental agency problem between managers and entrepreneurs, on the one hand, and capital markets, on the other hand. In this framework, wealth taxes distort the effort that managers and entrepreneurs expend to create firm value, which capital markets reward them for, by tilting their choice of projects towards less productive ventures. If wealth taxes are broad-based enough to affect individuals at the top of the wealth distribution, who contribute the most to firms' productivity, then they end up depressing capital accumulation and output in the economy. Our preliminary simulations show that even a simple version of the model accounts well for the degree of income and wealth inequality in the United States, including the distinct degrees of concentration of the distributions of income and wealth. The model implies a substantial aggregate output loss from wealth taxes of the magnitude currently being debated, which leads to a reduction in inequality that would be achieved at a much lower cost by modestly increasing the progressivity of income taxes rather than by introducing wealth taxes.



Home Bias and the Gender Wage Gap

Presenter: *Christian Hellwig*



Job Transformation, Specialization, and the Labor Market Effects of AI

Presenter: *Lukas Mann*

Coauthors: *Lukas B. Freund*

A central effect of automation is to transform jobs—shifting their task content. We develop a general-equilibrium model of this process. Occupations bundle tasks; workers possess task-specific skills and sort by comparative advantage. When a task is automated, remaining tasks gain in importance, so wage effects depend on workers' full skill profiles. We estimate the distribution of task-specific skills and project individual-level wage effects of generativeAI automation. Moderate exposure benefits workers on average but high exposure harms them, with large dispersion within occupations; the return to social skills rises, that to analytical skills falls; and low-earners gain more than high-earners. Job transformation drives these results.



Masculinity Norms, Occupational Choice, and Reallocation: Micro and Macro Evidence

Presenter: *Rachel Schuh*

Coauthors: *Martina Uccioli*

Over the past five decades, the labor force participation rate of American men has declined significantly—from 80% in the 1970s to just around 68% today—and similar trends are observed across the developed world. During the same period, the fastest-growing occupations have been those traditionally considered “feminine,” such as nursing. Are these two trends connected? Do masculinity norms discourage men from entering the labor force by deterring them from pursuing jobs perceived as feminine? We document that male employment participation declined more following sectoral reallocation in more gender-traditional labor markets, as proxied by occupational segregation by gender. We are then examining whether masculinity norms systematically distort men's occupational choices away from income-maximization and estimating men's willingness to pay (in terms of forgone earnings) to avoid occupations seen as feminine. Understanding the role of masculinity norms in shaping men's labor market decisions is crucial for addressing the ongoing decline in male labor force participation.



Self-employment as Self-insurance

Presenter: *Daniel Jaar*

I investigate the role played by microentrepreneurship as an income of last resort. Using microenterprise surveys from developing countries, I document that, on average, 19% of microentrepreneurs start firms because they do not find jobs. Relative to the rest, these ‘necessity microentrepreneurs’ are more likely to come from unemployment, run smaller and less profitable firms, and earn larger income gains when transitioning to wage employment, consistent with a self-insurance motive. To evaluate the channel’s aggregate implications, I propose a two-sector model of occupational choice with labor market frictions, where risk-averse workers separated from wage work can resort to informal microentrepreneurship to avoid unemployment. I find that enforcing entry regulations among microentrepreneurs may reduce welfare even if it increases output. Despite significant challenges due to limited state capacity, introducing unemployment insurance increases welfare and is strongly progressive.



The Changing Value of Employment

Presenter: *Davide Alonzo*

Coauthors: *Giovanni Gallipoli*

Are earnings informative about the returns from work? We estimate the value of worker-occupation matches in an equilibrium model that distinguishes between wages and latent returns and derive a measure of rents that has three properties: (i) it reflects wages and unobserved match values; (ii) it delivers a monetary metric for compensating differentials; (iii) it illustrates how the marginal workers within a match affect the rents of others. We show that earnings closely track rents between 1980 and 2018. The results highlight a dichotomy: while rents arise from latent values, their evolution is driven by technology.



Hiring with a Quantity-Skill Tradeoff

Presenter: *Lawrence Warren*

Coauthors: *Fil Babalievsky, Henry Hyatt,*



Artificial Intelligence and the Macroeconomy

April 24–25, 2026; Assembly Hall, Campus Center, University at Albany,
Uptown Campus, Albany, NY

Attendees:

Daniela Puzzello — *Indiana University*
Hao Pang — *University of Texas at Dallas*
Karan Bhasin — *University at Albany, SUNY*
Aakash Kalyani — *Federal Reserve Bank of St. Louis*
Andre Mouton — *Wake Forest University*
Gustavo de Souza — *Federal Reserve Bank of Chicago*

Andre Victor Luduvic — *Federal Reserve Bank of Cleveland*
Lukasz Drozd — *Federal Reserve Bank of Philadelphia*
Jaromir Nosal — *Boston College*
Ruben Fernandez Fuertes — *Bocconi University*
Marlon Azinovic-Yang — *University of North Carolina at Chapel Hill*
Niklas Bonnmann — *Vrije Universiteit Amsterdam*

Organizers:

Nick Pretnar — *LAEF, UCSB*
Ben Griffy — *University at Albany, SUNY*
Adrian Masters — *University at Albany, SUNY*



Simulating the Survey of Professional Forecasters

Presenter: Daniela Puzzello

Coauthors: Anne Lundgaard Hansen, John J. Horton, Sophia Kazinnik, and Ali Zarifhonarvar

We develop a framework for simulating professional expectations formation using large language models (LLMs). Combining novel hand-collected data on Survey of Professional Forecasters (SPF) participant characteristics with real-time macroeconomic data and lagged SPF median forecasts, we prompt LLMs to generate quarterly forecasts for 23 variables over 1999-2023. The resulting synthetic panel replicates key properties of the human survey, including forecast accuracy, median forecast levels, revision dynamics, and cross-sectional dispersion. Ablation exercises show that performance is largely driven by information conditioning: removing forecaster characteristics modestly worsens performance, whereas removing real-time data and especially lagged SPF medians leads to substantial error increases. Our framework offers a scalable complement to traditional surveys, enabling counterfactual analysis under alternative information sets, retrospective “as-of” forecasting, and rapid prototyping of survey designs.



Did I Make Myself Clear? The Fed and the Market under the 2020 Monetary Policy Framework

Presenter: Hao Pang

Coauthors: Anna Cieslak, Michael McMahon

We study the impact of the Federal Reserve’s communication on financial markets following the adoption of its revised policy framework in 2020. We propose a channel whereby market uncertainty stemming from perceived policy errors can raise risk premia. Post-framework communication introduced uncertainty about the Fed’s reaction function. Market concerns about policy mistakes amid incoming economic data drove up term premia, undermining easy financial conditions the Fed initially sought. While short-rate expectations were anchored by forward guidance, term premium sensitivity to inflation news increased. The subsequent shift in Fed’s words and then actions helped stabilize premia, mitigating adverse macroeconomic news.



Identifying Episodes of Fiscal Austerity: An LLM-Based Approach

Presenter: *Karan Bhasin*

Coauthors: *Prakash Loungani*

This paper introduces a hierarchical Large Language Model (LLM) framework for the automated identification of narrative fiscal shocks. We develop a multi-stage architecture to extract austerity episodes from IMF Article IV reports (2004–2020) for 17 OECD countries. Relative to manual coding, our approach improves replicability and auditability by generating a documented sequence of classification steps. Benchmarking against Adler et al. (2024), we find that the LLM-based classification aligns closely with the narrative benchmark, while differing on a small subset of episodes. Local projection estimates indicate that LLM-identified shocks are associated with smaller estimated multipliers than the narrative benchmark, with the difference linked in large part to differences in shock persistence and endogeneity.



Seniority Bias, Vintage Mismatch, and Productivity in the AI Era

Presenter: *Aakash Kalyani*



Measuring Task-Level Technological Exposure: A Language Model Approach

Presenter: *Andre Mouton*

This paper develops methods and tools for measuring the exposure of occupational tasks to technological substitution. Patent abstracts and task statements are matched and classified by a small, open-source language model. The model is fine-tuned and validated against a foundation AI, achieving accuracy improvements of roughly 5X over conventional ‘word embedding’ approaches. Model fine-tuning and a rules-based match threshold are critical for realizing these gains. The approach replicates stylized facts about IT exposure, but diverges sharply from survey-based measures of AI automation risk, which systematically understate exposure among high-wage occupations. A living dataset and Python package allow researchers to measure exposure across user-defined technology and task categories, with minimal time lag and at fine temporal resolution.



AI in the Office and the Factory: Evidence from Administrative Software Registry Data

Presenter: *Gustavo de Souza*

What is the effect of artificial intelligence (AI) on employment and wages? Using administrative data on a comprehensive set of AI applications developed in Brazil, I show that AI is widely used not only in administrative tasks but also in the production setting, where it primarily supports process optimization and quality control. Exploiting exogenous variation in AI development costs, I find that AI affects administrative and production workers differently. In office occupations, AI reduces employment, especially among young workers performing data-processing tasks. In production occupations, it increases employment of young and low-skilled workers by helping them perform decision-making tasks that previously required expertise. On net, AI increases employment in Brazil because gains in production outweigh losses in office work.



AI-Augmented Capital-Skill Complementarity

Presenter: *Andre Victor Luduvica*

Coauthors: *Roberto Pinheiro*



Understanding Growth Through Automation: The Neoclassical Perspective

Presenter: *Lukasz Drozd,*

Coauthors: *Mathieu Taschereau-Dumouchel, and Marina Mendes Tavares*

We study how advancements in automation technology affect the division of aggregate income between capital and labor in the context of long-run growth. Our analysis focuses on the fundamental trade-off between the labor-displacing effect of automation and its positive productivity effect in an elementary task-based framework featuring a schedule of automation prices across tasks linked to the state of technology. We obtain general conditions for the automation technology and technical change driving automation to be labor-share displacing. We identify a unique task technology that reconciles the Kaldor facts with the presence of automation along the balanced growth path. We show that this technology aggregates to the Cobb–Douglas production function—thus providing novel task-based microfoundations for this workhorse functional form. We employ our theory to study the connection between the recent declines in the labor share and the unique nature of the current, IT-powered wave of automation.



The Macroeconomic Implications of the Gen-AI Economy

Presenter: *Jaromir Nosal*

Coauthors: *Pablo Guerrón-Quintana, Tomoaki Mikami*

We study the potential impact of the generative artificial intelligence (Gen-AI) revolution on the US economy through the lens of a multi-sector model in which we explicitly model the role of Gen-AI services in customer base management. We carefully calibrate the input-output linkages and the size of the Gen-AI sector by using occupational task characteristics, and find large spillovers of the Gen-AI productivity gains into the aggregate economy. A 10% increase in productivity in the Gen-AI sector over a 10 year horizon implies a 5.6% increase in aggregate GDP, despite the AI sector representing only 20% of the overall economy. That shock also implies a significant reallocation of labor away from the Gen-AI sector and into non-AI sectors. We decompose these effects into parts coming from the input-output structure and customer base management and find that they each contribute roughly equally to the rise in GDP. In the absence of either channels, real GDP essentially does not respond to the increase in productivity in the Gen-AI sector.



Monetary Policy Shocks: A New Hope — Large Language Models and Central Bank Communication

Presenter: *Ruben Fernandez Fuertes*

I develop a multi-agent LLM framework that processes Federal Reserve communications to construct narrative monetary policy surprises. By analyzing Beige Books and Minutes released before each FOMC meeting, the system generates conditional expectations that yield less noisy surprises than market-based measures. These surprises produce theoretically consistent impulse responses where contractionary shocks generate persistent disinflationary effects, and enable profitable yield curve trading strategies that outperform alternatives. By directly extracting expectations rather than cleaning surprises ex post, this approach demonstrates how multi-agent LLMs can implement narrative identification at scale without contamination in high-frequency measures.



Deep Learning in the Sequence Space

Presenter: *Marlon Azinovic-Yang*

Coauthors: *Jan Žemlička*

We develop a deep learning algorithm for constructing globally accurate approximations to functional rational expectations equilibria of dynamic stochastic economies in the sequence space. We use deep neural networks to parameterize key equilibrium objects, such as policies or prices, as functions of truncated histories of exogenous shocks. We train the neural networks to satisfy equilibrium conditions along simulated paths of the economy. We illustrate the performance of our method in three environments: (i) a high-dimensional overlapping generations economy with multiple sources of aggregate risk; (ii) an economy with heterogeneous households and firms facing uninsurable idiosyncratic risk and large shocks to idiosyncratic and aggregate volatility; and (iii) a stochastic life-cycle economy with a continuous asset choice and a discrete early-retirement choice that induces local convexities in the continuation values of working-age cohorts. We also propose practical neural policy architectures that guarantee monotonicity of predicted policies, enabling the



A Global Solution Method for HACT Models with Aggregate Risk

Presenter: *Niklas Bonnmann*

We develop a method to compute global solutions to continuous-time heterogeneous agent (HACT) models with aggregate risk, non-stationarity, and financial constraints. When the cross-sectional distribution evolves stochastically, the associated master equation becomes second order. First, we provide a heuristic derivation of the second-order master equation with a stochastic Kolmogorov forward equation and show how its distributional derivatives can be approximated using a polynomial chaos expansion, thereby transforming it into a standard, albeit high-dimensional, partial differential equation (PDE). Second, building on the Deep BSDE approach for high-dimensional PDEs, we develop a general-equilibrium extension -- GE-Deep BSDE -- that jointly solves for value functions, equilibrium prices, and distributional consistency using multi-objective deep learning. Applying the method to a two-asset model, we show that while total factor productivity shocks fail to generate sufficient risk premia, stochastic capital depreciation with collateral constraints produces endogenous risk and helps reconcile microeconomic wealth dispersion with macro-finance evidence.



Demographics in the 21st Century

April 30–May 1, 2026; Federal Reserve Bank of Cleveland, Cleveland, OH

Attendees:

Raquel Fernandez — *New York University*
Agustin Diaz Casanueva — *Central Bank of Chile*
Lidia Cruces — *Goethe University, Frankfurt*
Ami Ko — *Georgetown University*
Stefania Albanesi — *University of Miami*
Irina Popova — *University of Bonn*
Minchul Yum — *Virginia Commonwealth University*
Mary Ann Bronson — *Georgetown University*
Fang Yang — *Federal Reserve Bank of Dallas*

Selahattin Imrohoroglu — *University of Southern California, Marshall School of Business*
Cheng Chen — *Clemson University*
Kathleen McKiernan — *Vanderbilt University*
Svetlana Pashchenko — *University of Georgia*
Toni Braun — *National Graduate Institute for Policy Studies (GRIPS), Tokyo*
Alexander Richter — *Federal Reserve Bank of Dallas*

Organizers:

Karen Kopecky — *Federal Reserve Bank of Cleveland*
Nick Pretnar — *LAEF, UCSB*
Ben Griffy — *University at Albany, SUNY*
John Jones — *Federal Reserve Bank of Richmond*



Understanding Latin America's Fertility Decline: Age, Education, and Cohort Dynamics

Presenter: *Raquel Fernández*

Coauthors: *Milagros Onofri, Inés Berniell, and Azul Menduiña*

This paper examines the sharp decline in fertility across Latin America using both period and cohort measures. Combining Vital Statistics, Census microdata, and UN population data, we decompose changes in fertility by age, education, and joint age–education groups. We show that the decline in period fertility between 2000 and 2022 is driven primarily by reductions in within-group birth rates rather than by changes in population composition, with the largest contributions coming from younger and less-educated women. Comparing the cohort born in the mid 1950s and the one born in the mid 1970s, we find that the decline in completed fertility reflects not only delayed childbearing but also substantial reductions in the average number of children per woman. This is driven primarily by lower fertility among mothers rather than by rising childlessness. Our findings provide new evidence on the nature of Latin America's transition to below-replacement fertility and highlight several open questions for future research.



The Effect of Cognitive Skills on Fertility Timing

Presenter: *Agustin Diaz Casanueva*

In the NLSY79, 69% of women in the lowest cognitive ability quartile have a first birth before age 22, compared with 22% in the highest. This gradient persists within education groups and within contraceptive method, suggesting that neither education sorting nor method selection fully accounts for it. I estimate a life-cycle model in which cognitive ability shifts the effectiveness of fertility control, beyond standard opportunity-cost channels. Restricting effectiveness to be equal across ability groups causes the model to underpredict the ability–fertility gradient by a factor of five. Equalizing effectiveness to high-ability levels reduces births before age 22 by 50% and raises college attendance by 15%. A cost-reduction policy generates welfare gains of 10% of lifetime consumption for the lowest ability quartile but near-zero gains for the highest, suggesting that improving effectiveness matters more than reducing cost for disadvantaged women.



Financial Incentives to Fertility: From Short to Long Run

Presenter: *Lidia Cruces*

Coauthors: *F. Javier Rodríguez-Román*

Are financial incentives effective in increasing fertility rates? Empirical evidence suggests they are, primarily in the short run (around implementation). Can such policies also increase the total number of children in the long run? We address this question by using a structural life-cycle model of fertility and labor supply, calibrated to replicate the short-run effects of a cash transfer paid at childbirth implemented in 2007 in Spain. The model incorporates labor market duality, a defining feature of Spanish labor markets that negatively impacts fertility. Our calibrated model replicates a 6% increase in fertility rates in the short run but only generates a 3% rise in completed fertility over women's lifetimes—the long run. Eliminating labor market duality increases lifetime fertility by 6.62%, but the discrepancy between short- and long-run effects of the incentive persists. These results highlight the limited impact of financial incentives alone to sustain fertility gains.



Childcare and Long-Term Care: Altruism and Strategic Motives within the Family

Presenter: *Ami Ko*

Coauthors: *Chao Fu, and Chamna Yoon*

In the U.S., families play a central role in the provision of childcare and long-term care (LTC), with grandparents frequently assisting their adult children with childrearing and middle-aged children providing much of the care for elderly parents. This paper studies the motives underlying these intergenerational time transfers and evaluates their policy implications. Using HRS data and exogenous variation in childcare costs, we first document evidence consistent with strategic implicit contracting, whereby early childcare assistance is repaid with later LTC. To quantify the underlying mechanisms and assess policy impacts, we estimate a dynamic model of an extended family who interact over care arrangements and bequests with both altruistic and strategic concerns. The estimated model shows that strategic reciprocal motives play a central role in shaping the intertemporal linkage between childcare and LTC. These mechanisms generate meaningful policy spillovers across care sectors. In particular, childcare subsidies that reduce grandparental involvement crowd out future family-provided LTC, while conditional expansions of old-age social insurance that reward past childcare provision reinforce the family's role in both care sectors.



Boomerang College Kids: Unemployment, Job Mismatch and Coresidence

Presenter: *Stefania Albanesi*

Coauthors: *Rania Gihleb, and Ning Zhang*

Labor market outcomes for young college graduates have deteriorated substantially in the last twenty five years, and more of them are residing with their parents. The unemployment rate at 23-27 year old for the 1996 college graduation cohort was 9%, whereas it rose to 12% for the 2013 graduation cohort. While only 25% of the 1996 cohort lived with their parents, 31% for the 2013 cohort chose this option. Our hypothesis is that the declining availability of ‘matched jobs’ that require a college degree is a key factor behind these developments. Using a structurally estimated model of child-parent decisions, in which coresidence improves college graduates’ quality of job matches, we find that lower matched job arrival rates explain two thirds of the rise in unemployment and coresidence between the 2013 and 1996 graduation cohorts. Rising wage dispersion is also important for the increase in unemployment, while declining parental income, rising student loan balances and higher rental costs only play a marginal role.



The Demographic Cliff and the Market for Higher Education: Implications for Public Finance

Presenter: *Irina Popova*

We construct a general equilibrium life-cycle model of the college market with heterogeneous colleges, student college quality and major choice, and subsequent labor market outcomes of workers in different occupations. The cross-sectional distribution of college quality and major specialization is an equilibrium outcome and is shaped by the demographic structure of the economy, by public education spending and college loan policies, as well as by the endogenous relative wages that college graduates with different majors command in the labor market. We use the model to evaluate the aggregate and distributional consequences of the “demographic cliff”, a scenario with lower population growth that will reduce the number of high-school graduates in the coming decades. We then discuss how public education finance reforms by the current administration shape these outcomes.



Parental Leave Policies, Fertility, and Labor Supply

Presenter: *Minchul Yum*

Coauthors: *Daisoon Kim*

Parental leave policies are often viewed by policymakers as a potential tool for addressing low fertility and gender inequality. We develop a dynamic heterogeneous-household life-cycle model in which couples jointly choose careers, labor supply, savings, childcare, fertility, and parental leave use. The model incorporates a segmented labor market where career-oriented jobs are difficult to re-enter after interruptions and is calibrated to Korean data. Our quantitative results show that more generous parental leave benefits can simultaneously increase fertility and narrow gender gaps in labor supply and wages over the life cycle. These effects arise through a career-retention channel: job-protected leave allows more women to remain attached to career-oriented jobs during childrearing years, improving long-run career progression despite temporary labor-supply reductions. Without job protection or labor-market segmentation, these effects weaken substantially or reverse, highlighting the importance of broader labor-market institutions in shaping parental leave policy effects.



The Power of Policy Incentives: Female Labor Supply, Fertility and Parental Leave Policy Design

Presenter: *Mary Ann Bronson*

Parental leave (PL) policies vary widely across countries, creating different incentives for childbearing and maternal labor supply. Yet, most empirical studies find that PL policy has negligible effects on these variables. This paper provides new evidence about the effects of PL policy design on women's outcomes. We first show that the standard quasi-experiment in the literature, which relies on regression discontinuities around PL reform dates, understates the true impact of PL policies for a simple reason: treatment and control groups share the same policy exposure for most of the evaluation period, making it difficult to isolate meaningful differences in outcomes. Using event studies from the U.K., Sweden, Austria, and Germany, we then show that women's labor force participation, PL take-up, and fertility choices closely track the incentives embedded in each country's policy design. Next, we build a dynamic model of women's work and fertility decisions, allowing us to simulate the effects of policies with varying durations, compensation structures, and eligibility criteria. The model replicates the small quasi-experimental estimates in the literature but shows that the true policy effects are far larger, in some cases increasing women's employment by 10 percentage points over the ten years after first birth, while also boosting fertility. Calibrating the model to U.K. data, we show that when exposing U.K. women to Swedish, Austrian, or German policies, they make work and childbearing choices that mirror those in those countries, confirming the power of policy incentives. Finally, we examine policy designs that maximize both labor force participation and fertility. We find that policies with moderate leave duration and generous wage replacement simultaneously boost employment and birth rates. Our findings indicate that the design of parental leave matters far more than past estimates imply, with lasting effects on women's careers and family choices.



Marriage and Work Among Prime-Age Men

Presenter: Fang Yang

Coauthors: Adam Blandin, John Bailey Jones

Married men work more hours than men who have never been married. Fixed effect regressions reveal that part of this gap is attributable to an increase in work around the time of marriage. Two potential explanations for the increase are: (i) men hit by positive labor market shocks are more likely to marry; and (ii) marriage leads men to work more hours. Using a structural life-cycle model, we find that marriage substantially increases male hours of work. Counterfactual simulations suggest that declining marriage rates account for roughly half of the fall in prime age male hours observed over recent decades.



Optimal Pension Reforms during Demographic Transitions

Presenter: Selahattin Imrohoroglu

This paper investigates the long-term implications of demographic aging for public pension sustainability, using population projections from both international and national sources. Drawing on the 2024–2100 United Nations World Population Prospects (UN WPP) as well as forecasts from the Social Security Administration (SSA), U.S. Census Bureau (CB), Congressional Budget Office (CBO), and Japan's National Institute of Population and Social Security Research (IPSS), I document substantial cross-country variation in projected old-age dependency ratios. In addition, there is a large range of forecasts produced by different assumptions on fertility and immigration. I calculate a key forward-looking metric—the number of full retirement age (FRA) postponements required to maintain the 2023 dependency ratio through 2100. I show that advanced economies like Japan, having aged much earlier, face more modest required adjustments, while many middle-income countries will need much sharper policy shifts. I further demonstrate the sensitivity of unfunded pension liabilities to alternative demographic projections in the U.S. and Japan, and quantify the impact of Japan's 2004 macroeconomic slide reform on long-run fiscal sustainability. The results underscore the importance of proactive pension reform—particularly in developing economies entering demographic transition with limited time and institutional and fiscal capacity.



Population Aging, Firms, and Worker Composition

Presenter: *Fang Yang*

Coauthors: *Adam Blandin, John Bailey Jones*

Population aging changes the composition of workers in the labor force. We develop a framework to understand how worker composition affects productivity in an economy with heterogeneous firms and apply it to Japan. Reallocating workers toward the optimal composition of experience and skill would increase output by 3.8%. Including worker aging and reallocation costs, we find strong complementarity between workers' skill and firm productivity, which generates differential hiring strategies across firms and heterogeneous effects of population aging on wages and market concentration. Using Japanese establishment data, we validate key model predictions: larger firms disproportionately hire young workers, and regions with slower working-age population growth



The Elasticity of Social Security Policy

Presenter: *Kathleen McKiernan*

Coauthors: *Neha Bairoliya*



Redistribution, Distortions, and the Welfare Effects of Social Security

Presenter: *Svetlana Pashchenko*

Coauthors: *Youngsoo Jang, and Ponpoje Porapakkarm*

What is the best way to reform Social Security? Academic literature offers diverging advice. There is a well-known result that the optimal size of Social Security is zero, implying it is best to phase the program out. Other studies argue that much can be gained by redesigning the program, given its current size. We provide a unified analysis that examines how the optimal size of Social Security depends on the key features of its design. We first develop a theoretical decomposition tracing the program's welfare effects to (i) income redistribution, (ii) distortions on the annuitization level, and (iii) intertemporal distortions. We then quantitatively assess the role of these channels. We show that the zero-optimal-size result arises because Social Security is too distortive and not redistributive enough. Once these design flaws are corrected, it is even optimal to increase the size of the program.



Why Aging Induces Deflation and Secular Stagnation

Presenter: *R. Anton Braun*

Coauthors: *Daisuke Ikeda*

We provide a quantitative theory of deflation and secular stagnation. In our lifecycle framework an aging population puts persistent downward pressure on the price level, real interest rates, and output. A novel feature of our theory is that it also recognizes the reactions of government policy. The central bank responds to falling prices by reducing its policy nominal interest rate and the fiscal authority responds by allowing the public debt–GDP ratio to rise.



The Postpandemic US Immigration Surge: New Facts and Inflationary Implications

Presenter: *Alexander W. Richter*

Coauthors: *Anton Cheremukhin, Sewon Hur, Ronald Mau, Karel Mertens, and Xiaoqing Zhou*

The U.S. experienced an extraordinary postpandemic surge in unauthorized immigration. This paper combines administrative data on border encounters and immigration court records with household survey data to document two new facts about these immigrants: They tend to be hand-to-mouth consumers and low-skilled workers that complement the existing workforce. We build these features into a model with capital, household heterogeneity, and population growth to study the inflationary effects of this episode. Contrary to the popular view, we find little effect on inflation, as the increase in supply was largely offset by an increase in demand.



LAEF

Laboratory for Aggregate Economics and Finance

Director – *Finn Kydland*

Assistant Director – *Nick Pretnar*

Design and Production – *Thill Creative*

UC **SANTA BARBARA**

Laboratory for Aggregate
Economics and Finance

LAEF

Laboratory for Aggregate Economics and Finance

2112 North Hall

University of California, Santa Barbara

Santa Barbara, CA 93106–9215

U.S.A.

econ-laef@ucsb.edu

www.laef.ucsb.edu