Research Statement

Joachim Hubmer
University of Pennsylvania

July 2026

I am a quantitative macroeconomist studying how economic growth and technological change shape the distribution of income and wealth. Growth has distributional consequences because it changes both what the economy produces and how it produces it. In this statement, I organize my research contributions around two linked processes. The first concerns how production generates income, how that income is divided between labor and capital, and how economic activity is distributed across firms. The second concerns how households save and convert income into wealth over the life cycle and across generations.

On the production side, I show why labor's share of U.S. national income, roughly stable for decades, has declined since 1980. Two opposing forces shape it: rising household income shifts spending toward labor-intensive goods and services, while cheaper and more capable equipment and software induce firms to substitute capital for workers. After 1980, the technological force dominated, accounting for most of the decline. In related work, I find that fixed adoption costs lead large firms to adopt new automation technologies first. Together, this technological change and its uneven adoption explain why the aggregate labor share falls even though the typical firm's labor share barely changes, and why firm sales concentration rises. I also show that firms' investments in reaching customers can raise concentration without necessarily making markets less competitive.

On the household side, I show how business returns, saving, inheritance, and tax policy shape wealth concentration. Using population-wide Norwegian administrative data spanning more than two decades, I show that inherited wealth and within-generation accumulation through unusually high business returns and saving rates are both central at the top. My work on the United States finds that the decline in tax progressivity was the dominant force behind the post-1980 rise in wealth concentration.

Across these projects, I combine detailed microeconomic data with structural quantitative models. I estimate or calibrate the models' parameters using data ranging from household expenditure surveys and population registers to firm censuses and administrative records. The models trace how changes in technology and policy interact with household behavior, firm decisions, and market frictions to shape aggregate and distributional outcomes.

1. Technology, Production, and Factor Income Distribution

My research in this area includes:

[1] Hubmer, J. (2023): "The Race Between Preferences and Technology", *Econometrica*, 91, 227–261.

[2] Hubmer, J. and P. Restrepo (2026): "Not a Typical Firm: Capital–Labor Substitution and Firms' Labor Shares", *American Economic Journal: Macroeconomics*, 18(2), 34–71.

[3] Hubmer, J., M. Chan, S. Ozkan, S. Salgado, and G. Hong (2026): "Scalable versus Productive Technologies", PIER Working Paper.

[4] Hubmer, J. and L. Nord (2026): "Investment in Demand and Dynamic Competition for Customers", PIER Working Paper.

[5] Abraham, A., P. Gottardi, J. Hubmer, and L. Mayr (2023): "Tax Wedges, Financial Frictions and Misallocation", *Journal of Public Economics*, 227.

One branch of my research studies how technology and firm heterogeneity shape the distribution of income across workers, capital owners, and firms. The distributional effects of technological progress depend on where it occurs, which firms adopt it, and how markets reallocate customers, workers, and capital across firms. My work uses this perspective to study the determinants of the falling labor share, rising firm concentration, growing markups of prices over production costs, and misallocation (the failure of labor and capital to flow to the firms where they would be most productive). A common theme that emerges from my research is that these patterns are not isolated phenomena: they are connected outcomes of technological change, firm heterogeneity, and market frictions. Throughout, I combine aggregate data with micro data on households and firms, from expenditure surveys to census and administrative firm records, and I use quantitative models matched to these data to evaluate the distributional and policy consequences of these forces.

The falling labor share. In Hubmer (2023), I ask why the U.S. labor share of national income declined after 1980. Its earlier stability had long been a classic empirical regularity of economic growth. Although the analysis focuses on the postwar U.S. economy, the broader question is when successive waves of technological change embodied in capital goods, from mechanization through computers and software to current automation and AI, lower labor's share.

I find that two opposing forces shape the aggregate labor share. On the demand side, I show empirically that richer households spend more on labor-intensive goods and services. To demonstrate this pattern, I link household spending data from the Consumer Expenditure Survey to measures of the labor and capital used along each good or service's value chain, constructed from detailed BEA input-output tables. With stable non-homothetic preferences, this fact implies that income growth tends to push the labor share up. On the production side, I document that labor-share declines are concentrated in the production of goods and services whose full value chains use equipment and software capital intensively, precisely the forms of capital whose prices have fallen most rapidly, so capital-embodied technical change tends to push down the labor share for those goods and services. I combine these empirical facts in a multi-sector growth model with non-homothetic preferences and capital–labor substitution, estimating the key preference and production elasticities from the same micro data. The result is a unified account of both the pre-1980 stability and the later decline of the labor share: for decades the two forces roughly offset each other, but after 1980 the technological force became dominant, accounting for roughly 80% of the decline between 1982 and 2012. The broader lesson is that both the direction and pace of technological progress matter: growth does not automatically reduce labor's share, but when advances are concentrated in capital technologies that substitute for workers, gains can shift persistently toward capital owners.

In Hubmer and Restrepo (2026), we ask where this technological force appears within industries. The key micro fact is that in industries where labor shares fell sharply, such as manufacturing, the median firm's labor share was essentially unchanged between 1982 and 2012. This rules out a uniform shift toward more capital-intensive production across all firms and points instead to changes concentrated in specific firms.

This pattern is often interpreted as evidence that market power has risen among "superstar" firms. We find instead that adoption of new capital-intensive technologies is highly uneven. In Census Bureau surveys of automation use, large firms adopt much more; and in firm-level production functions estimated from Compustat data, large firms show the strongest increases in the elasticity of output with respect to capital. The evidence therefore points to technological change that is biased not only toward capital, but also toward large firms.

New automation technologies often involve fixed costs: firms must reorganize production, install new systems, or automate entire tasks. Large firms adopt first because they can spread these fixed costs over more sales. In a task-based model, production consists of many distinct tasks (assembly, bookkeeping, quality control), each performed by either workers or machines; automation means that new capital goods take over tasks previously performed by workers. A firm-dynamics model with this task structure and fixed costs of automation matches the firm-level evidence and shows that capital–labor substitution can account for most of the aggregate labor-share decline, even though the median firm's labor share barely moves. The task structure matters because capital advances can reduce labor shares even when substitution across existing production tasks is limited: if new capital technologies expand the set of tasks performed by capital, they create an automation margin beyond the usual substitution margin. This helps reconcile micro estimates of limited substitution, such as Oberfield and Raval (2021), with aggregate estimates that imply stronger substitution, such as Karabarbounis and Neiman (2014).

The same mechanism also changes the market structure. Automation lowers labor shares especially at large, high-markup firms and reallocates sales toward them. In the model, this raises concentration and generates a modest increase in the aggregate markup through a composition effect. These changes in market structure therefore arise from the same technological change that lowers the labor share.

Productivity, scalability, and firm size. Why are some firms much larger than others within the same narrowly defined industry? Standard firm-dynamics models, following Lucas (1978) and Hopenhayn (1992), usually answer this question with total factor productivity (TFP): firms with high TFP produce more output from a given bundle of inputs, while all firms share a common technology for scaling up production. In Hubmer, Chan, Ozkan, Salgado, and Hong (2026), we study a different and complementary margin: scalability. We measure scalability through returns to scale, the elasticity of output with respect to all inputs. A firm can be productive, scalable, or both. Conceptually, TFP is a level advantage: how much a firm produces from a given input bundle. Scalability is a scale advantage: how much output can expand as input use grows, and how quickly diminishing returns set in.

This distinction creates two empirical challenges. First, productivity and scalability must be separately measured, even though both can make firms large. Second, high returns to scale at large firms could reflect either movement along a common non-homothetic production function, with large firms operating where returns to scale happen to be high, or ex ante differences in technologies across firms. We distinguish these possibilities using firm panel data from thirteen countries, combining administrative tax records that cover the universe of incorporated firms in Canada, plant-level U.S. Census manufacturing data, and harmonized financial statements for eleven European countries (Orbis). We estimate flexible production functions nonparametrically and use an iterative clustering procedure to separate movement along common production functions from differences across production functions.

We find that returns to scale rise strongly with firm size, especially at the top of the size distribution, while TFP flattens or declines among the largest firms. The evidence points mostly to ex ante differences in technologies rather than movement along a common production function. Among the largest firms, size is therefore more closely associated with scalability than with higher TFP.

These findings matter for macroeconomic misallocation. Statically, a given distortion, such as a tax wedge or a financing constraint, produces larger output losses at highly scalable firms: because their returns diminish slowly as they expand, such firms can be left operating far below their efficient size. Dynamically, financial frictions slow scalable firms' growth toward that efficient size, and they also distort who chooses to become an entrepreneur, because they change the value of operating a scalable firm. Quantitatively, allowing firms to differ in scalability in an otherwise standard entrepreneurship model more than doubles the estimated efficiency costs of financial frictions, relative to a conventional calibration with heterogeneity only in TFP.

Demand investment and competition. Firms compete not only by choosing prices and technologies, but also by investing to shape demand. They spend resources on sales, marketing, and customer acquisition so that customers know about them and buy from them. In Hubmer and Nord (2026), we study how this demand investment shapes competition, allocative efficiency, and product-market policy. The spending is large in Capital IQ data, about two-thirds the size of physical investment among publicly traded U.S. firms, and directly affects which buyer-seller relationships form and persist.

In the model, customers see only a subset of suppliers (a search friction), and firms can invest to reach more potential customers. When several suppliers reach the same customer, they compete for that customer's demand. The central innovation is that competition itself becomes an equilibrium outcome: firms' demand investment determines how many suppliers compete for each customer, and therefore how competitive the market is. A firm's price and markup depend not only on its own fundamentals, but also on which competitors reach the same customers and how those competitors price. A second distinctive feature is that prices are allocative on two margins: they affect sales per customer, and they affect whether customers form, retain, or switch supplier relationships. This structure sets the model apart from existing models of customer acquisition, and it lets the model match salient facts on customer and sales dynamics, markups, and spending on sales and marketing.

In equilibrium, firms' demand investment has two opposing effects. It draws customers away from rivals; this is privately valuable to the firm, but socially it is a form of business-stealing, since one firm's gain is another's loss. It also benefits consumers by creating buyer-seller matches that would otherwise not form. In the calibrated model, the business-stealing force dominates, so firms over-invest: the social planner would reduce aggregate demand investment by about one third. But demand investment is also pro-competitive: reducing it too much can leave customers with fewer suppliers and raise markups. We therefore find that a tax on demand investment works best when combined with a production subsidy, which offsets the market-power distortion that the tax alone would amplify.

The paper also changes how we interpret market concentration. If large firms invest more in reaching customers, demand investment can raise concentration without necessarily raising market power. Concentration is therefore not a sufficient statistic for competition: by itself, it cannot show how competitive a market is. More generally, the model distinguishes what creates markups from what makes differences in markups costly. Search frictions determine how many suppliers compete for each customer and thereby generate markups. By contrast, the cost of markup differences depends on how strongly customers shift purchases toward lower-priced suppliers. Standard models use a single elasticity for both channels. In our model, variation in sales per customer identifies the second response separately. Because it is small, the observed markup differences imply a substantially lower efficiency cost than in standard models.

Taxation and production reallocation. A related, more theoretical contribution studies what firm heterogeneity implies for corporate-tax incidence, meaning who ultimately bears the tax. In Abraham, Gottardi, Hubmer, and Mayr (2023), we revisit the classical Harberger (1962) result that, in a closed economy, the burden of corporate taxation falls almost entirely on capital. We show that this result is fragile once entrepreneurs differ in wealth and ability and choose their occupation, organizational form, investment, and financing structure. Financial frictions and the tax advantage of debt over equity generate misallocation of capital and talent; a corporate tax increase then reallocates production from C-corporations to pass-through firms, lowering labor productivity and wages. Corporate-tax incidence therefore depends on the organizational and financial margins through which production is reallocated, and workers can bear a substantial share of the burden.

Work in progress. My ongoing work uses linked administrative firm data to study how firms build production networks. In joint work with Salvatore Lo Bello, Claudio Michelacci, and Lukas Nord, we use Italian business-to-business transaction data to decompose firm growth into gains from adding customers and suppliers and gains from larger transactions with existing partners, and we use identified innovation events to estimate how innovation reshapes these networks. Innovating firms expand sales and expenditures mainly by adding customers and suppliers. We embed these facts in a model in which firms choose customer and supplier relationships while setting prices, extending the demand-side framework of Hubmer and Nord (2026) to the supplier side. We use the framework to assess how these relationships shape competition, aggregate productivity, and spillovers from innovation.

2. Inequality and Wealth Accumulation

My research in this area includes:

[1] Hubmer, J., E. Halvorsen, S. Ozkan, and S. Salgado (2026): "Why Are the Wealthiest So Wealthy? New Longitudinal Empirical Evidence and Implications for Theories of Wealth Inequality", forthcoming at *Econometrica*.

[2] Hubmer, J., P. Krusell, and A. A. Smith, Jr. (2021): "Sources of U.S. Wealth Inequality: Past, Present, and Future", *NBER Macroeconomics Annual*, 35, 391–455.

[3] Hubmer, J. (2018): "The Job Ladder and Its Implications for Earnings Risk", *Review of Economic Dynamics*, 29, 172–194.

A longstanding question in macroeconomics is why the distribution of wealth is far more concentrated than the distribution of labor earnings. In benchmark incomplete-markets models such as Bewley (1986), İmrohoroğlu (1989), Huggett (1993), and Aiyagari (1994), households cannot fully insure against earnings shocks, so precautionary saving is the main source of wealth differences. This question is connected to the production-side themes discussed in Section 1, because top wealth is closely tied to business ownership. Many households at the top of the wealth distribution become wealthy not only by saving out of labor income, but by owning and operating successful firms. My work in this area asks which additional forces, including returns on business wealth, saving behavior, inheritances, and tax policy, account for both wealth differences at a point in time and the dynamics of wealth accumulation.

Drivers of wealth inequality. In Hubmer, Halvorsen, Ozkan, and Salgado (2026), we use Norwegian administrative data to answer directly a basic question that previous work could only tackle indirectly: how did today's wealthiest households accumulate their fortunes over the life cycle? The contribution is to use households' observed life-cycle paths to distinguish which theories of top wealth are most consistent with the empirical evidence. Previous work could often match cross-sectional wealth concentration. But inherited wealth, differences in investment returns or patience, saving motives that vary with wealth, and a small chance of exceptionally high earnings can each generate that concentration while implying different consequences for inheritance taxes and taxes on capital income and wealth.

We assemble Norwegian administrative panel data from 1993 to 2019 covering the entire population and the full household budget constraint: wealth, labor and capital income, taxes and transfers, and inheritances and inter vivos transfers. Norway is uniquely suited for this question: most of these components are third-party reported to the tax authorities rather than self-reported, and a shareholder registry links households to the firms they own. Because the data follow the same households for 26 years, the analysis captures not only levels of these components but also their timing and compounding over the life cycle. We then develop an intertemporal budget-constraint decomposition that attributes wealth gaps between any two groups to differences in saving rates, inheritances, returns on net wealth, and labor income. Applying it to the top of the wealth distribution, where direct dynamic evidence has been especially scarce, we find that for households aged 45–64 the excess wealth of the top 0.1% relative to the median is accounted for by higher saving rates (36%), higher inheritances (31%), and higher returns (28%), with higher labor income contributing only 5%.

The decomposition reveals a sharp distinction within the top. Households we label "New Money" start below median wealth and rise to the very top through exceptionally high business returns, high saving rates, and high labor income. Households we label "Old Money" start very rich, mostly because of inheritances and inter vivos transfers, and continue to accumulate through above-average saving rates and moderately above-average returns; labor income contributes little to their excess wealth.

We then use these life-cycle decompositions to test the implications of different theories of top wealth. Five standard models, each emphasizing a familiar mechanism such as earnings risk, heterogeneous returns, discount-factor heterogeneity, non-homothetic bequests, or exceptionally high labor income, can match cross-sectional patterns of labor income, total income, and wealth concentration. But they fail to jointly generate the dynamics of New and Old Money, with most placing too much weight on inherited wealth and too little on within-generation accumulation through returns and saving. The model based on exceptionally high labor income has the opposite problem. A model with heterogeneous entrepreneurial returns and decreasing returns to scale, combined with non-homothetic preferences over consumption and saving, fits the dynamic evidence much better; this mechanism ties top wealth directly to the kind of firm heterogeneity studied in Section 1. Theories of top wealth should therefore be judged not only by whether they match who is rich at a point in time, but also by whether they match how households become and remain rich over the life cycle.

Whereas the previous paper uses direct life-cycle evidence to study how households reach the top, in Hubmer, Krusell, and Smith (2021) we ask what drove the long-run evolution of U.S. wealth inequality. U.S. top wealth shares fell through much of the twentieth century and then rose sharply after 1980. We study this evolution with a benchmark incomplete-markets model featuring stochastic returns whose mean and dispersion vary with wealth, capturing the different portfolios that rich and poor households hold.

We calibrate the model to reproduce the U.S. wealth distribution in 1967, the start of our sample period, and then supply it with the observed historical paths of its driving forces: declining tax progressivity from published effective-tax-rate series, rising top earnings and greater earnings risk as estimated in survey and tax data, changing factor shares, and asset returns by asset class. These forces operate jointly through equilibrium saving, wages, interest rates, and the different portfolios of rich and poor households, placing the leading explanations for rising U.S. wealth inequality within a single quantitative environment. The framework shows not only which forces move top wealth shares in the right direction, but also why some prominent explanations are quantitatively limited once they are embedded in a heterogeneous-agent equilibrium model.

Our key finding is that the decline in tax progressivity since the late 1970s is the dominant force behind the rise in top wealth shares. Labor-income changes play a more nuanced role. A thicker upper tail of earnings raises top wealth shares modestly, but the broader rise in earnings risk induces precautionary saving below the top and can compress wealth concentration. Changes in factor shares have limited effects on top wealth concentration in the model, while portfolio differences across wealth groups and movements in asset returns are essential for understanding the short- and medium-run dynamics, especially the U-shaped path of top wealth shares. If tax progressivity remains at its current historically low level, the model implies further substantial increases in top wealth shares over the coming decades.

Earnings risk. The nature of earnings risk provides a bridge between firm heterogeneity and household wealth accumulation. In Hubmer (2018), I ask whether the non-Gaussian earnings dynamics documented in administrative data can be generated by a structural labor-market model, rather than imposed as a reduced-form statistical process. The empirical motivation is that earnings growth is not well described by the log-normal processes commonly used in macroeconomic models: Guvenen, Karahan, Ozkan, and Song (2021) and related work document large negative skewness and excess kurtosis in earnings changes. In plain terms, most workers experience small changes in a typical year, a few experience very large ones, and the largest changes are more often losses than gains.

I show that a life-cycle job ladder model with heterogeneous firms, on-the-job search, savings, and human-capital depreciation in unemployment can reproduce these facts. The model is calibrated to aggregate labor-market transition rates, the distribution of earnings in levels, and the average earnings loss after job displacement; no moments of earnings changes are used in the calibration. The model nonetheless reproduces the negative skewness and excess kurtosis measured in the administrative data. These patterns arise because workers climb gradually toward better firms, while job loss can push them far down the ladder; recovery then depends on search incentives, wealth, and skill loss. The model connects measured earnings risk to economic mechanisms: part of observed earnings dynamics reflects workers' movement across firms and their search and saving responses rather than primitive exogenous shocks. Counterfactuals that shut down workers' search and saving choices shrink the higher-order moments of earnings changes by 30–90%, depending on the moment and horizon. This provides a structural interpretation of earnings risk, a key input into models of consumption, saving, and wealth accumulation.

Work in progress. With Elin Halvorsen, Serdar Ozkan, and Sergio Salgado, I am pursuing two additional projects that connect top wealth dynamics more directly to policy and to firm technology. The first reverses the lens of Hubmer, Halvorsen, Ozkan, and Salgado (2026): instead of selecting households that end up at the top and tracing them backward, we select households that begin with substantial inherited wealth and follow them forward. Using the same intertemporal budget-constraint approach, we decompose the fate of inherited fortunes: how much households preserve and expand by investing successfully in businesses, how much they consume, how much they pass on, and how much they dissipate through low returns and failed firms. This provides direct dynamic evidence on what inherited wealth is used for, a key input for evaluating inheritance taxation. Preliminary results show high aggregate persistence at the top of the wealth distribution alongside a substantial subgroup of "squanderers" who run down most of their inherited wealth, a subgroup larger than standard models predict. We plan to combine these forward-looking profiles with estimates of the short-run responses of saving and bequests to historical variation in Norwegian inheritance tax rates, and to use the structural model to extrapolate to medium-run effects. The ultimate goal is a quantitative assessment of inheritance taxation: who gains, who loses, and at what efficiency cost.

The second project applies the firm-level estimation approach of Hubmer, Chan, Ozkan, Salgado, and Hong (2026) to Norwegian firm data and links the resulting estimates of TFP and returns to scale to the household histories of Hubmer, Halvorsen, Ozkan, and Salgado (2026). Because the Norwegian register data cover both firms and the households that own them, we can match top wealth holders to the technologies of the firms they operate. We ask whether New Money and Old Money own firms with different productivity and returns to scale, and whether scalability itself is partly an endogenous outcome rather than a fixed technological primitive. The answer matters for capital-income and wealth taxation: if top wealth holders operate high-returns-to-scale firms, a given tax wedge can generate larger efficiency losses; if returns to scale are endogenous, taxation also distorts technology choice and firm growth.

3. Summary and Outlook

Taken together, my work studies how the economy generates income and who ends up owning the wealth it creates. Looking ahead, my ongoing projects extend each branch, from the anatomy of production networks on the firm side to the uses of inherited wealth on the household side, and increasingly connect the two: the Norwegian data make it possible to link the technologies firms use to the wealth accumulated by the households that own them. My agenda thus remains centered on the question that runs through all of my work: when growth changes production, who receives and accumulates the gains?

References

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

Bewley, T. F. (1986): "Stationary Monetary Equilibrium with a Continuum of Independently Fluctuating Consumers," in *Contributions to Mathematical Economics in Honor of Gerard Debreu*, ed. by W. Hildenbrand and A. Mas-Colell, Amsterdam: North-Holland, 79–102.

Guvenen, F., F. Karahan, S. Ozkan, and J. Song (2021): "What Do Data on Millions of U.S. Workers Reveal About Lifecycle Earnings Dynamics?" *Econometrica*, 89(5), 2303–2339.

Harberger, A. C. (1962): "The Incidence of the Corporation Income Tax," *Journal of Political Economy*, 70(3), 215–240.

Hopenhayn, H. A. (1992): "Entry, Exit, and Firm Dynamics in Long Run Equilibrium," *Econometrica*, 60(5), 1127–1150.

Huggett, M. (1993): "The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies," *Journal of Economic Dynamics and Control*, 17(5–6), 953–969.

İmrohoroğlu, A. (1989): "Cost of Business Cycles with Indivisibilities and Liquidity Constraints," *Journal of Political Economy*, 97(6), 1364–1383.

Karabarbounis, L. and B. Neiman (2014): "The Global Decline of the Labor Share," *Quarterly Journal of Economics*, 129(1), 61–103.

Lucas, R. E. Jr. (1978): "On the Size Distribution of Business Firms," *Bell Journal of Economics*, 9(2), 508–523.

Oberfield, E. and D. Raval (2021): "Micro Data and Macro Technology," *Econometrica*, 89(2), 703–732.