VOLUME I • ISSUE VIII

Economic Research Briefs

THE UNIVERSITY OF CHICAGO
BECKER FRIEDMAN INSTITUTE
FOR ECONOMICS

bfi.uchicago.edu

The latest commentary and analysis from BFI's leading scholars in shorter, digestible forms.

Contents

VIEW ALL BRIEFS

bit.ly/4h0YM3M

RESEARCH BRIEF • OCTOBER 2025

Artificial Writing and Automated Detection

Based on BFI Working Paper No. 2025-116, "Artificial Writing and Automated Detection," by Brian Jabarian and Alex Imas, University of Chicago

Commercial AI detection tools significantly outperform open-source alternatives, with Pangram achieving near-zero error rates while open source options misclassify up to 78% of human text as AI-generated. A new policy framework allows institutions to systematically compare detectors based on their tolerance for false accusations versus missed AI content.

Generative Artificial Intelligence tools have been adopted faster than any other technology on record, giving rise to writing that is either assisted or entirely completed by Large Language Models (LLMs). The ubiquity of AI-generated writing across domains such as school assignments and consumer reviews presents a new challenge to stakeholders aiming to detect whether content was written by humans. While automated detection tools hold promise, their accuracy claims are difficult to verify since they rely on proprietary data and methods.

In this paper, the authors audit the set of leading AI detection tools and offer a framework to evaluate how they should be incorporated into potential policies. They evaluate four detectors—three commercial tools (Pangram, OriginalityAI, GPTZero) and one open-source baseline (RoBERTa)—on their ability to minimize two critical errors: false positives (wrongly flagging human text as AI) and false negatives (missing actual AI content).

Understanding detector performance requires grasping how these tools work. A detector evaluates text and assigns it a score, such that higher scores imply a greater likelihood that the text is AI-generated. A detector's performance

Figure 1 · Detector False Positive Rates by Genre

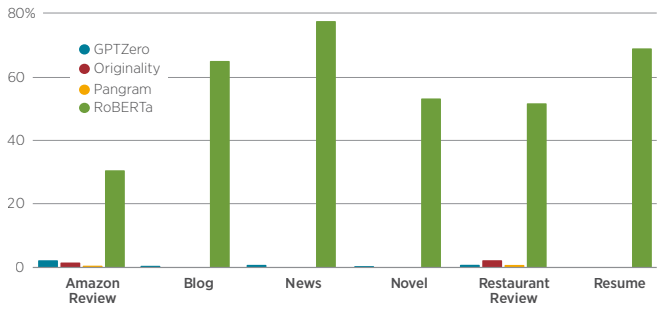

depends critically on the score threshold at which it classifies content as AI-generated. A higher threshold implies that the detector requires a higher score to classify a passage as AI-generated; this will naturally decrease the false positive rate while at the same time increasing the false negative rate.

Given this threshold trade-off, the authors evaluate performance in two ways: First using detector-optimized thresholds calculated to maximize the difference between true positive rate and false positive rate, and second by manipulating thresholds exogenously to demonstrate how policy designers can adjust detector settings based on their tolerance for different types of errors.

For their evaluation, the authors use a 1,992-passage text corpus that spans six everyday genres (news, blogs, consumer reviews, novels, restaurant reviews, and résumés) as input for their evaluation. Verified human-generated text is matched with AI-generated text using four frontier LLMs (GPT-4.1, Claude Opus 4, Claude Sonnet 4, Gemini 2.0 Flash). They also examine the effectiveness of AI "humanizers" (StealthGPT) in potentially bypassing detectors.

They find the following:

- Commercial detectors significantly outperform open-source alternatives across all metrics and AI models. Among commercial options, Pangram achieves essentially zero false positive rates and false negative rates on medium-length to long passages, both when using detector-optimized thresholds and exogenously-set thresholds. The false positive rate and false negative rate increase slightly on short passages, but remain well below reasonable policy thresholds.
- The performance gap between commercial and open-source tools is substantial. OriginalityAI and GPTZero constitute a secondary tier among commercial detectors with partial strengths, making the choice between the two dependent on the user's priority: minimizing false positive rate favors GPTZero, while maximizing ability to distinguish AI from human text favors OriginalityAI. In contrast, the open-source RoBERTa base is deemed unsuitable for high-stakes applications, misclassifying most human text with false positive rates of approximately 30-78% across scenarios.
- Pangram's false negative rate is robust to the use of current "humanizers" and remains low even when AI-generated passages are modified using tools such as StealthGPT. The other detectors are less robust to humanizers, with GPTZero largely losing its capacity to detect AI-generated text, showing false negative rate scores around 50% and above across most genres and LLM models.
- After converting vendor fees into cost per correctly flagged AI passage, Pangram is two times cheaper than OriginalityAI and is almost three times cheaper than GPTZero both overall and on shorter passages. Cost-per-true-positive analysis sharpens the price gap, making Pangram the most cost-efficient detector.
- The policy caps framework, which sets exogenous thresholds to test detector robustness, reveals that Pangram is the only detector that meets stringent policy requirements without compromising the ability to detect AI text. When policy caps are set at very low levels (0.5% false positive rate), Pangram continues detecting AI content effectively while other detectors see sharp degradation in their detection capabilities.

It is important to note that the implications of AI detection for writing and text-based work more generally are not obvious. LLMs are incredibly valuable tools that can facilitate idea generation and help tighten writing. At the same time, the use of LLMs to off-load a task where the receiver explicitly desires human input creates a host of agency problems. The use of AI text detectors in practice must thus strike a delicate balance to avoid discouraging the former while mitigating the issues posed by the latter.

READ THE WORKING PAPER

NO. 2025-116 · AUGUST 2025

Artificial Writing and Automated Detection

bfi.uchicago.edu/working-papers/artificial-writing-and-automated-detection

ABOUT OUR SCHOLARS

Brian Jabarian
Howard and Nancy Marks Fellow and Roman Family Center for Decision Research Principal Researcher, Chicago Booth

Alex Imas
Roger L. and Rachel M. Goetz Professor of Behavioral Science, Economics and Applied AI and Vasilou Faculty Scholar, Chicago Booth

THE UNIVERSITY OF CHICAGO
BECKER FRIEDMAN INSTITUTE FOR ECONOMICS

Written by Abby Hiller • Designed by Maia Rabenold

RESEARCH BRIEF • OCTOBER 2025

Firms Have Partial Knowledge: Evidence from a Reform

Based on BFI Working Paper No. 2025-120, "Firms Have Partial Knowledge: Evidence From a Reform," by Avner Strulov-Shlain, University of Chicago

A pricing reform in Israeli supermarkets that forced firms to adjust prices revealed that they operated with only partial knowledge of optimal pricing. When forced to explore new pricing strategies, firms gradually learned and improved performance, though meaningful deviations from optimal pricing persisted.

Economists typically assume that firms are sophisticated and make optimal decisions. Yet, in practice, firms often make mistakes. This study asks why: Do firms *know* the right action but fail to implement, or do they lack knowledge of the optimal action?

To explore this question, Strulov-Shlain examines a 2014 Israeli reform that banned prices ending in denominations of coins that no longer existed, in this case the equivalents of the U.S. penny and nickel. The change forced supermarkets to round prices to the nearest 10 Agorot (e.g., 2.90 or 3.00).

Importantly, the reform eliminated the possibility of "99" endings, which are optimal due to left-digit bias (the tendency for consumers to overweight the leftmost digit in a price). In this way, the reform created a natural experiment: Could firms predict the new optimal price endings?

Strulov-Shlain uses data covering supermarket prices and transactions from 2013 to 2015 to examine prices before and after the reform. He also measures how "crossing the left-digit threshold," e.g., changing a price from $4.99 to $5.00, affects demand.

He finds the following:

- Before the reform, roughly 45% of supermarket prices ended in 99, the optimal ending under left-digit bias, while nearly none ended in 00. This suggests firms were partially knowledgeable about optimal pricing.
- When prices ending in 99 were banned, about 40% of prices moved to 90-endings (the new optimal format), but 20% shifted to 00-endings, crossing the left-digit threshold. This suggests that firms lacked full knowledge of how left-digit bias would translate under new pricing.
- Crossing the left-digit threshold was costly. Changing a product's price from 4.99 to 5.00 reduced demand by 5-9% compared

Figure 1 · Shares of Price Endings Over Time

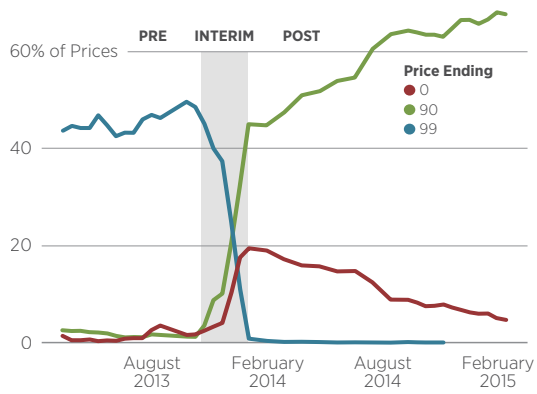

Note: The figure plots the shares of 99-, 90-, and 00-ending prices across products and stores by sampling period. The shaded area represents the time between announcement of the reform on October 17, 2013, and its enactment on January 1, 2014. The blue line represents 99-ending prices, the green line 90-ending prices and the red line 00-ending prices.

to changing it to 4.90. Despite this sizable effect, many firms initially adopted suboptimal 00-endings, suggesting that mispricing was due to lack of knowledge, not inability to implement.

- The share of 00-ending prices fell from roughly 20% immediately after the reform to 5-8% within a year, and around 3% by 2018–2019. Some retail chains exhibited abrupt, "light-bulb" shifts rather than gradual adjustment. At the product level, once prices moved to 90-endings, they tended to stay there, while those initially set to 00 rapidly reverted to 90. Products that transitioned from 00 to 90 were 15 times more likely to remain at 90 than revert, suggesting incremental learning through experience.
- To quantify the implications of these behavioral patterns, Strulov-Shlain builds a model that translates the estimated demand responses into optimal pricing prediction. He finds that firms lost about 0.75 percentage points in profit immediately after the reform due to mispricing but recovered within a year as they adapted. By the end of the adjustment period, average pricing efficiency exceeded pre-reform levels, suggesting genuine learning rather than mechanical adaptation.

These findings challenge the standard economic assumption that firms effectively optimize, suggesting instead that partial knowledge should be treated as a benchmark for firm behavior. The mechanism sustaining partial knowledge appears to be insufficient exploration. Had firms experimented more extensively with different price endings before the reform, they would have learned the full structure of demand. For researchers, these findings encourage taking partial optimization seriously when modeling firm behavior. Assuming full optimization may lead to incorrect predictions about how firms respond to regulatory changes or market shocks.

READ THE WORKING PAPER

NO. 2025-120 · SEPTEMBER 2025

Firms Have Partial Knowledge: Evidence From a Reform

bfi.uchicago.edu/working-papers/firms-have-partial-knowledge-evidence-from-a-reform

ABOUT OUR SCHOLAR

Avner Strulov-Shlain
Associate Professor of Marketing, Chicago Booth

CHICAGO BOOTH
THE UNIVERSITY OF CHICAGO BOOTH SCHOOL OF BUSINESS

Written by Abby Hiller • Designed by Maia Rabenold

RESEARCH BRIEF • OCTOBER 2025

Human Capital Accumulation Across Space

Based on BFI Working Paper No. 2025-134, "Human Capital Accumulation Across Space," by Klaus Desmet, Southern Methodist University; Dávid Krisztián Nagy, Centre de Recerca en Economia Internacional; and Esteban Rossi-Hansberg, University of Chicago

Regions with lower education costs maintain persistently higher levels of human capital and development over centuries. Improving educational access in poor regions generates local benefits but may reduce global welfare when population shifts away from more productive areas. Policies that equalize education costs within regions can produce unintended consequences as people relocate toward less productive locations, offsetting the benefits of cheaper schooling.

Human capital, as measured by levels of school-based education, is unevenly distributed across space. In 2000, people in the Netherlands had an average of 10.8 years of schooling, compared to 2.5 years in the Central African Republic. Comparing inhabitants of the most and least educated corners of the globe—specifically, 1° × 1° grid cells at the 90th and the 10th percentiles of educational attainment—this range goes from 11.8 to 3.4 years.

What drives these large differences in human capital across space? The authors examine how two factors shape the geography of development, both today and in the future. First, the cost of acquiring human capital varies widely across locations—in some places, access to education is relatively expensive or difficult, limiting the supply of human capital. Second, the productivity of human capital differs across locations—where human capital generates higher returns, demand for

human capital: the collective skills, knowledge, and abilities of individuals that can be used to create economic value

Figure 1 · Years of Schooling and Cost of Education

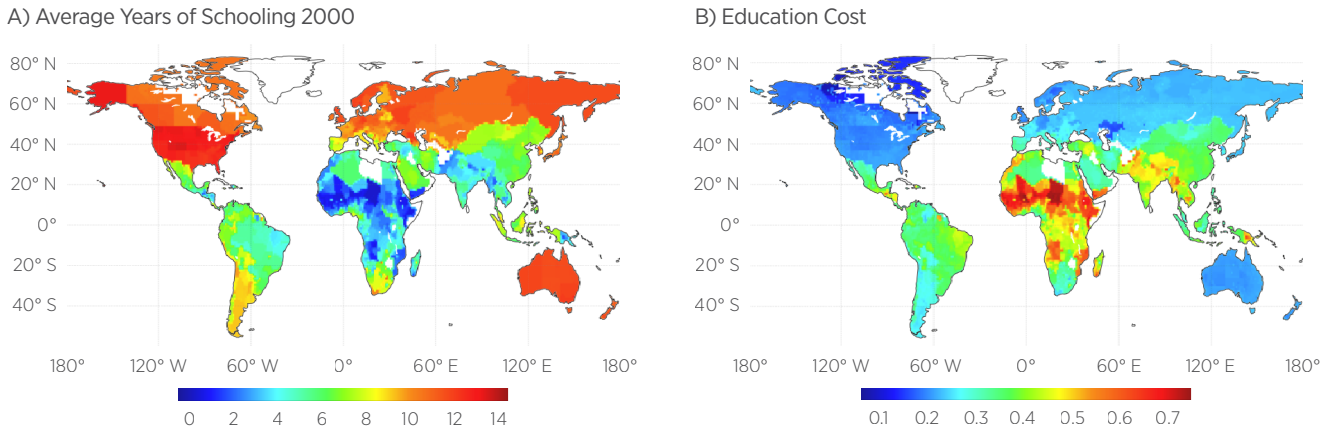

Note: Panel A shows the average years of schooling in year 2000. Panel B shows the education cost, as recovered from the model.

Figure 2 · Cost of Education and GDP per Capita

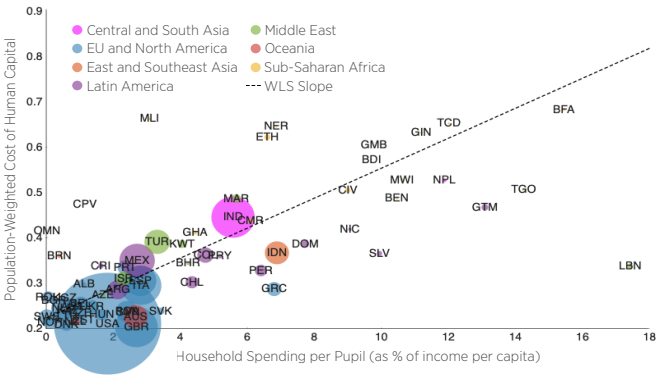

Note: This figure displays the relationship between the population-weighted cost of human capital and household spending per pupil as a share of income per capita.

it will be greater. These forces interact with migration, trade, and innovation to determine how human capital evolves across the globe over time.

To address this question, the authors develop a dynamic spatial model of the world economy at a 1° × 1° resolution. In the model, individuals choose where to live and how much human capital to acquire, taking into account that both moving between locations and upgrading human capital are costly, and these costs vary across space. Firms in each location produce differentiated goods using labor, human capital, and land, with trade between locations subject to transport costs.

The model incorporates two key productivity forces. First, a location's overall productivity benefits from **agglomeration economies**—as population density increases, so does productivity. Second, locations accumulate human-capital-augmenting technology over time through two channels: local innovation (which depends on the local stock of human capital) and diffusion from other locations. This creates dynamic feedback loops where human capital today boosts innovation tomorrow, attracting more people and generating further innovation.

The authors quantify the model using data on population, income, and schooling from 2000. They identify location-specific costs of acquiring human capital by matching the model to observed changes in schooling levels between 2000 and subsequent years. They then simulate the model forward for 200 years to project how human capital and economic development evolve across space.

They find the following:

- The model predicts strong persistence in the geography of development. Over the span of two centuries, today's developed regions, such as coastal Australia, Western Europe, Japan, and the United States, remain the most developed 200 years from now. The same persistence holds for population density, as locations that are dense today continue to be dense two centuries from now.
- Even after 200 years, the world economy remains far from reaching a steady state where all regions grow at the same rate. This finding stands in sharp contrast to spatial models that ignore human capital, which instead predict that poor but densely populated areas will eventually catch up to wealthier regions through agglomeration effects alone.
- Low education costs drive persistent advantages in developed regions. Because highly developed regions tend to have low education costs, their human capital levels tend to be high, both in the short and the long run. This advantage is magnified by dynamic feedback loops over the transition path. *Current* levels of human capital improve *future* productivity, because human capital is an input in the growth of human-capital-augmenting productivity.
- The model reveals a strong negative correlation between education costs and local economic fundamentals, as places with better amenities, higher productivity, and more favorable conditions for development also tend to have cheaper access to education. As a result, the low cost of acquiring human capital in the developed world keeps these regions ahead, generating the persistence the model projects.
- Reducing education costs by the same percentage across poor regions generates local gains but may cause global losses. The authors examine counterfactual policies that lower the cost of human capital acquisition while maintaining the relative differences between locations within a region. Whether implemented in sub-Saharan Africa, Latin America, or Central and South Asia, the local economy benefits: human capital levels rise, both in the short

agglomeration economies: the productivity benefits that arise when economic activity concentrates in a particular location, as firms and workers gain advantages from being near one another

and the long run. This enhances innovation in human-capital-augmenting productivity, generating positive dynamic effects.

- Higher local welfare retains a larger share of the global population in the target region, further reinforcing productivity through agglomeration economies. However, the effects on global welfare differ markedly by region. When the policy is implemented in a low-income region, like sub-Saharan Africa and Central and South Asia, the local increase in population comes at the expense of regions with better economic fundamentals, reducing global agglomeration and innovation. In contrast, when implemented in a middle-income region like Latin America, the population reallocation comes partly from regions with worse fundamentals, improving outcomes globally.
- Equalizing educational costs across space may lead to unintended consequences. As an alternative policy, the authors consider setting the cost of education to the same level across all grid cells within a region, such as sub-Saharan Africa—eliminating the variation in education costs between locations. Because schooling costs are higher in less developed areas within the region, equalizing costs lowers them more dramatically in those areas.
- This creates incentives for population to migrate toward locations with weaker economic fundamentals. As a result, an increasingly larger share of the region's population resides in less productive locations, which hurts overall innovation and weakens agglomeration economies through greater geographic dispersion. These negative effects may partly or even fully offset the positive impact of cheaper access to human capital.

These findings suggest that effective development policies must account for spatial frictions, agglomeration effects, and the dynamic relationship between human capital and productivity across space. When evaluating education policies, policymakers must consider not only local benefits but also how these policies reshape the global distribution of population and economic activity. Policies that retain population in regions with weak economic fundamentals may generate local gains while producing global losses. Moreover, within-region heterogeneity matters: equalizing access to education across locations with varying economic potential can trigger population movements that undermine the policy's intended benefits.

READ THE WORKING PAPER

NO. 2025-134 · OCTOBER 2025

Human Capital Accumulation Across Space

bfi.uchicago.edu/working-papers/human-capital-accumulation-across-space

ABOUT OUR SCHOLAR

Esteban Rossi-Hansberg
Glen A. Lloyd Distinguished Service Professor, Kenneth C. Griffin Department of Economics
THE UNIVERSITY OF CHICAGO | KENNETH C. GRIFFIN DEPARTMENT OF ECONOMICS

Written by Abby Hiller • Designed by Maia Rabenold

RESEARCH BRIEF • OCTOBER 2025

Jealousy of Trade: Exclusionary Preferences and Economic Nationalism

Based on BFI Working Paper No. 2025-138, "Jealousy of Trade: Exclusionary Preferences and Economic Nationalism,"
by Alex Imas, University of Chicago; Kristóf Madarász, London School of Economics; and Heather Sarsons, University of Chicago

Many voters support tariffs and protectionist policies that materially hurt them because they derive value from consuming or possessing goods that others want but do not have. Individuals with such "exclusionary preferences" are significantly more willing to accept higher prices from tariffs than from other policies, like stimulus spending.

Why do voters support protectionist policies that materially harm them? Recent evidence shows that tariffs raise consumer prices and generate retaliatory trade measures that reduce employment. Such policies remain politically popular, however, particularly among those most economically affected. In this paper, the authors offer a new explanation: support for nationalist economic policies stems from a fundamental desire for dominance, which generates preferences for excluding others from consumption opportunities.

The authors build on prior research documenting that a substantial portion of the population derives utility not just from consuming goods, but from consuming goods that others desire but cannot obtain. They incorporate this desire for dominance into a model of international trade, showing that such exclusionary preferences reduce the value of trade and generate support for restrictive policies. The model predicts that people with exclusionary preferences will support tariffs that harm both their own consumption and their trading partner's consumption, but will show no such preference for policies that affect only domestic consumption.

To test these predictions, the researchers conduct two surveys. They begin by measuring respondents' exclusionary preferences using an incentivized experimental method in which participants bid on a unique good under three scenarios with varying degrees of exclusion of other potential buyers. Those whose willingness to pay increased with the level of exclusion are classified as having "preferences for exclusion," a pattern observed in roughly 40% of respondents (consistent with prior research). Respondents are then randomly assigned to evaluate tariff policies under different conditions and asked about their support for various economic policies.

The authors find the following:

- Exclusionary preferences strongly predict tariff support, but only when tariffs harm trading partners. Those with exclusionary preferences are 12.3 percentage points more likely to support a 15% tariff that would raise prices domestically. When respondents are told the tariff would not harm the foreign country, support between those with and without exclusionary preferences is statistically indistinguishable.

- Those with exclusionary preferences are more accepting of inflation caused by tariffs than by other policies. When comparing support for tariffs versus stimulus policies that would generate identical 15% price increases, respondents with exclusionary preferences show significantly higher support for tariffs.
- Exclusionary preferences predict support for a broad range of protectionist policies that harm domestic consumers. Beyond tariffs, those with exclusionary preferences are significantly more likely to support policies explicitly designed to maintain consumption gaps between nations, even when informed these policies would raise prices for Americans. They also show higher support for restricting foreign investment, emphasizing that the US should "come out on top" in trade relations, and limiting purchases from foreign countries. These patterns held across different trading partners (China, Mexico, and Canada), suggesting the effects are not driven by hostility toward specific nations.
- The relationship between exclusionary preferences and policy support is not explained by political ideology or cognitive biases. While political preferences partially mediate the relationship (Democrats are less likely to hold exclusionary preferences), the core association remains strong and statistically significant after controlling for party affiliation and zero-sum thinking (a cognitive bias where people believe gains for some come at others' expense).

These findings have important implications for understanding the political economy of trade policy. The results suggest that voter support for protectionist measures may be driven less by misunderstanding of economic costs or by narrow self-interest than by a fundamental preference for policies that exclude foreign consumers from consumption opportunities, even at personal economic cost. This helps explain why tariffs remain politically popular despite clear evidence that they raise prices and harm employment. The findings also suggest that inflation stemming from protectionist policies may generate less political backlash than equivalent price increases from other sources, as voters with exclusionary preferences view such costs as more acceptable when they serve to limit foreign consumption.

READ THE WORKING PAPER

NO. 2025-138 · OCTOBER 2025

Jealousy of Trade: Exclusionary Preferences and Economic Nationalism

bfi.uchicago.edu/working-papers/jealousy-of-trade-exclusionary-preferences-and-economic-nationalism

ABOUT OUR SCHOLARS

Alex Imas
Roger L. and Rachel M. Goetz Professor of Behavioral Science, Economics and Applied AI and Vasilou Faculty Scholar, Chicago Booth

Heather Sarsons
Associate Professor of Economics and William Ladany Faculty Scholar, Chicago Booth

CHICAGO BOOTH
THE UNIVERSITY OF CHICAGO BOOTH SCHOOL OF BUSINESS

Written by Abby Hiller • Designed by Maia Rabenold

RESEARCH BRIEF • OCTOBER 2025

Superstar Firms Through the Generations

Based on BFI Working Paper No. 2025-118, "Superstar Firms Through the Generations," by Yueran Ma, University of Chicago; Benjamin Pugsley, University of Notre Dame; Haomin Qin, London Business School; and Kaspar Zimmermann, University of Hamburg

New technologies that exhibit economies of scale, that confer low adoption costs for new entrants, and that require organizational learning, give rise to superstar firms for a long period of time. These firms enjoy systematic advantages relative to both firms that came before and potential entrants thereafter.

In recent years, researchers and the business media have focused attention on superstar firms, or that small set of top companies that account for a large share of output. However, though we understand a good deal about the current make-up of these firms, questions persist about how these firms are born, and whether/how these firms attain their superstar status over time.

In this work, the authors collect new data to conduct an extensive analysis of the largest US companies over the past century. These data include the 2018 Fortune list (a recent example year before COVID), which covers the largest 1,000 companies by sales across all sectors; the first Fortune list in 1955, which covers the largest 500 industrial (i.e., manufacturing and mining)

Figure 1 · Birth Years of 2018 Fortune 1,000 Companies

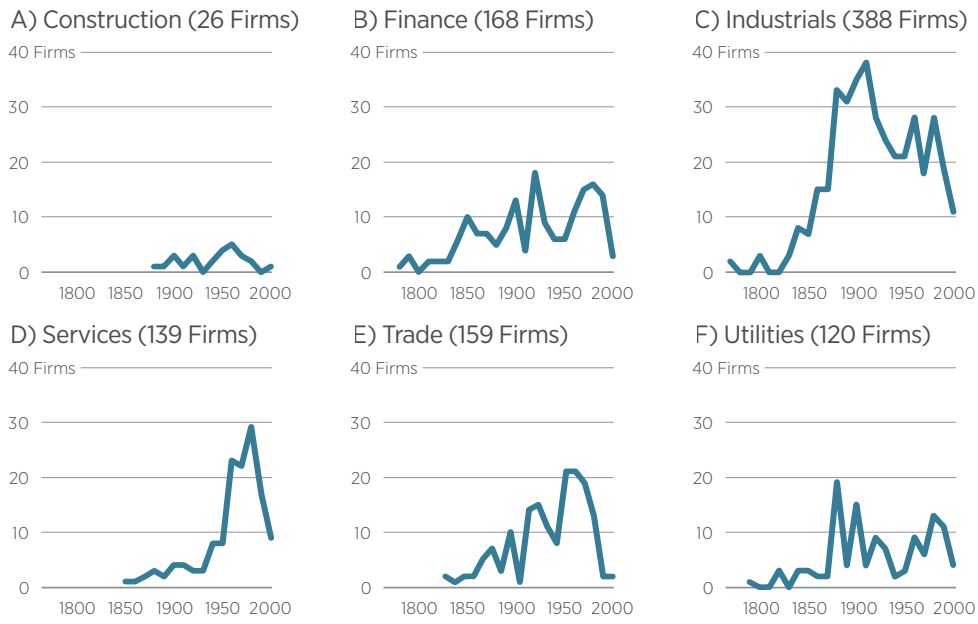

Note: This figure plots the number of 2018 Fortune 1,000 companies that originated in each decade, based on the firm's industry in 2018 (the main sectors correspond to Standard Industrial Classification codes; please see working paper for more details). For industrials, there is substantial clustering of birth years around the turn of the 20th century: the cohort from 1880 to 1920 is the most represented among today's largest industrials.

Figure 2 · Age Distribution of the Largest Industrials

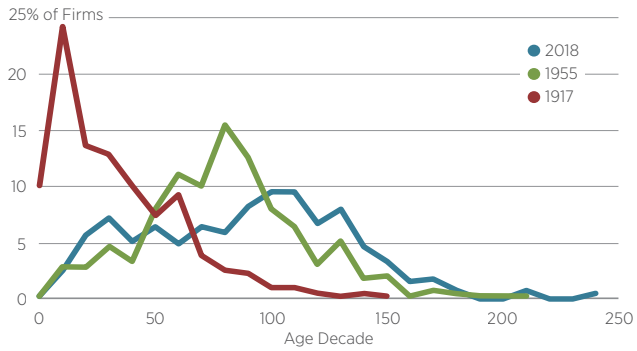

Note: The blue line shows the age decade distribution of the largest 388 industrials (by sales) in the 2018 Fortune list (which contains 388 industrial firms). The green line shows the age decade distribution of the largest 388 industrials (by sales) in the 1955 Fortune list. The red line shows the age decade distribution of the largest 388 industrials (by assets) in 1917.

companies by sales; Fortune lists of top retailers and wholesalers; and a list of the largest 500 industrials by assets in 1917. The authors also research the origin story for each firm using extant resources.

For reference, the authors use "superstar firms" to denote large companies that have achieved an extraordinary size relative to other businesses (for example, total sales of the top 1,000 firms in the 2018 Fortune list exceed 40% of U.S. private sector gross output). A review of the data reveals the following:

Industrials

- The emergence of superstar **industrial firms** in the United States is highly uneven over time—with a special generation from 1880 to 1920 (around the time of the **Second Industrial Revolution**) remaining dominant among the largest manufacturers in the 1910s, 1950s, and 2010s—but top firms still experience substantial turnover. This "special generation" has a lasting influence; the median age of top firms was around 30 in 1917, 60 in 1955, and 100 by 2018.
- However, the persistence of this special generation does not extend to individual firms. Twenty-one percent of the top industrials on the 1955 Fortune list remain on the 2018 Fortune list. Correspondingly, among the 388 industrials in the 2018 Fortune list, 137 were born between 1880 and 1920, but only 50 were among the top 388 industrials in 1955; the rest (and the majority) is represented by "late bloomers."
- By comparison, the authors find a similar pattern in the industrial history of Germany. In the UK, though, superstar firms are relatively young today. In part, this evidence suggests that the special generation among top industrials is not just a result of country-specific regulations (given the similarities in the US and Germany), or military buildup in the world wars (which would also be relevant in the UK).

Retail and wholesale

- In contrast to large industrial firms, superstar retailers and wholesalers are relatively stable in their age distribution. Today, the birth years of large retailers and wholesalers cluster around 1960 to 1980. In the 1950s, though, top companies primarily date back to around 1900.
- The authors show that the largest retailers and wholesalers have stayed 60 to 70 years old on average, without a special generation.

Service

- Today, birth years for superstar service firms cluster around 1960 to 1980, and few large ones existed before then; also, few services companies would qualify for the largest businesses in the economy until recently.
- In 2018, the large services in Fortune 1,000 are young, with a median age of 43, which resembles the youth of top industrials in 1917.
- A possible explanation for the relative youth of service superstars is that the cohort of services firms born around the **Third Industrial Revolution** forms a special generation, like the cohort of industrial firms born around the Second Industrial Revolution, but several more decades of data are necessary to confirm this account.

industrial firms: firms in manufacturing and mining industries according to the Standard Industrial Classification codes, which are four-digit codes that classify a company by its economic activity

Second Industrial Revolution: a period from around 1860 to 1900, marked by advancements in steel, electricity, and petroleum, and giving rise to large-scale industrialization and major corporations

Third Industrial Revolution: begun in the late 20th century, this era is characterized by the shift from mechanical and analogue technologies to digital electronics; key drivers include computers, the internet, and telecommunications

The authors take these data to a model of firm dynamics, which offers the following explanations for the birth and persistence of superstar firms:

- Declining adoption costs of the scalable technology generate the advantage of the special cohort relative to firms that came before;
- The accumulation of productivity over time through learning gives superstars first-mover advantages relative to potential entrants afterwards;
- Idiosyncratic firm-level shocks keep the top firms changing despite the persistence of the special generation; and
- The persistence of the special generation does not necessarily imply staleness or lack of dynamism among top firms.

Bottom line: Certain historical settings produce special generations of entrants that give rise to superstar firms for decades to come. These settings occur occasionally, and include new technologies that exhibit economies of scale, that confer low adoption costs for new entrants, and that require organizational learning. The combination of these forces produces special cohorts that have a strong edge relative to both firms that came before and potential entrants thereafter. That said, individual superstars keep churning due to idiosyncratic shocks.

Figure 3 · Age Distribution of the Largest Retailers and Wholesalers

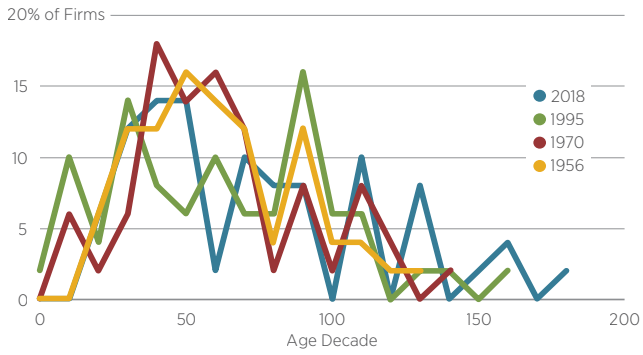

Note: The blue line shows the age decade distribution of the largest 50 retailers and wholesalers in the 2018 Fortune list. The green line shows the age decade distribution of the largest 50 in the 1995 Fortune list. The red line shows the age decade distribution of the largest 50 in 1970. The yellow line shows the age decade distribution of the largest 50 in 1956.

READ THE WORKING PAPER

NO. 2025-118 · AUGUST 2025

Superstar Firms Through the Generations

bfi.uchicago.edu/working-papers/superstar-firms-through-the-generations

ABOUT OUR SCHOLAR

Yueran Ma
Carhart Family Professor of Finance, Chicago Booth

CHICAGO BOOTH
THE UNIVERSITY OF CHICAGO BOOTH SCHOOL OF BUSINESS

Written by David Fettig • Designed by Maia Rabenold

RESEARCH BRIEF • OCTOBER 2025

The Breakdown of the English Society of Orders: The Role of the Industrial Revolution

Based on BFI Working Paper No. 2025-108, "The Breakdown of the English Society of Orders: The Role of the Industrial Revolution," by Cara Ebert, RWI Leibniz Institute for Economic Research; Leander Heldring, Northwestern University; James A. Robinson, University of Chicago; and Sebastian Vollmer, University of Goettingen

The Industrial Revolution fractured the strictures from a centuries-old "society of orders" in England, introducing the phenomenon of social mobility that created new opportunities for workers, and initiating a "dawn of liberty" for those previously entrenched in rigid social hierarchies.

Few areas of economic research have attracted as much attention in recent decades as issues surrounding social mobility, both among researchers and policymakers. Even as economists argued about causation (does moving to a new neighborhood, for example, reduce intergenerational poverty, or does selection bias explain positive outcomes?), policymakers initiated programs to help Americans move.

However, social mobility has long been a feature of western economies, though the drivers of that mobility have changed over time. This is especially true of the dramatic impact that the **Industrial Revolution** had on the traditional "**society of orders**," which divinely ordered allocations of people into separate parts of a

Figure 1 · Relationship of 17th Century Probate Inventory Wealth and Number of Hearths

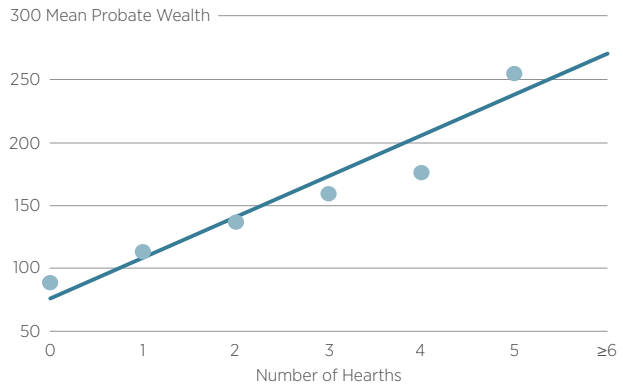

Note: This figure shows a binned scatter plot (mean probate inventory wealth for each categorical number of hearths) and a linear fitted line of the probate inventory wealth on the number of hearths. The sample consists of 269 matched hearth tax observations to probated wealth from a dataset of probates that cover the hearth tax period. Please see working paper for more details.

Industrial Revolution: a period of rapid technological advancement beginning in Great Britain in the mid-18th century, that transformed economies from agricultural and handicraft-based systems to those dominated by machine manufacturing and large-scale factories

society of orders: the system of social organization in pre-industrial Europe that divided society into a rigid hierarchy of status groups, where an individual's position was determined primarily by birth and fixed by law and custom

hierarchy, and which existed in English society in the centuries leading to the transformative 18th century and beyond. The authors of this new work are the first to empirically conceptualize the society of orders, which allows them to examine its consequences for social mobility over time.

Prior to the Industrial Revolution, English society was shaped by a rigid social hierarchy where people's wealth and status were largely predetermined by birth, noble titles, and family surnames. Medieval society was composed of three orders, those who work, pray, and fight, with surnames often indicating specific occupations that sons would inherit from their fathers (Smiths were blacksmiths, for example, Millers ground wheat, Bakers baked bread, and Coopers forged barrels). Also, surnames largely indicated, if not determined, societal status (children born into such surnames had no chance to rise above their station). That this "great chain of being" was ordained by God ensured that people kept their place, at least until the onset of an unexpected economic "miracle" broke those chains.

To investigate the impact of the Industrial Revolution on English society, the authors assembled two unique historical datasets: their pre-Industrial Revolution data consists of hearth tax returns from 1662-1674, which recorded household wealth based on the number of fireplaces, covering 343,022 heads of households across 26 counties, while their post-Industrial Revolution data includes digitized Principal Probate Registry records from 1862-1899, documenting the wealth of deceased persons who held at least 5 pounds at death. Both datasets include individuals' names, titles, and locations, allowing the researchers to analyze how well traditional social markers predicted wealth distribution across these two periods. They find the following:

- Before the Industrial Revolution, being a member of the **nobility** or **gentry** explains 17% of the variation in the share of wealth individuals owned, afterwards it explains only 6%, a decline of two-thirds.
- Family surnames, which historically carried strong occupational and status associations, explained 10% of wealth variation in the pre-industrial period but only 6% afterward.
- This breakdown was not uniform across England but was most pronounced in regions that experienced greater industrialization. In heavily industrialized areas, the explanatory power of surnames for wealth distribution fell to essentially zero, while in less industrialized regions, surname-based wealth prediction remained constant at around 10%.

The authors conducted benchmarking exercises to confirm that these data do, indeed, represent a breakdown of traditional society and not just statistical noise. They reveal that before the Industrial Revolution, the existing social hierarchy captured about 41% of its theoretical maximum explanatory power for wealth distribution, suggesting a society where traditional orders still significantly influenced economic outcomes. After the Industrial Revolution, this figure plummets to less than 10% of the theoretical maximum.

The authors also investigate the extent to which different attributes predict whether a person will be wealthy, or in the top decile of the wealth distribution. Essentially, the Industrial Revolution has little impact on the probability that a noble is rich, while the probability that a member of the gentry is wealthy increases significantly. The authors interpret this effect on the gentry as reflecting the flexibility of this tier of the society of orders. As the economy grew, new entrants were naturally wealthier than incumbents who were less able to take advantage of new economic opportunities.

This research also addresses two questions relating to geographic mobility and social mobility: Where did people with socially mobile surnames tend to end up geographically and within what economic

noble: member of the highest-ranking aristocrats, whose titles were often inherited and whose ranks, in descending order, were duke, marquis, earl, viscount, and baron

gentry: a class below noble that included landowners, knights, baronets, and esquires, often living on rental income from their estates and often held local positions like justices of the peace

sectors; and what were the characteristics of socially mobile **parishes**? They find that:

- People with socially mobile surnames, or those showing increased wealth dispersion within family lines, were significantly more likely to migrate to northern England and find employment in manufacturing industries.
- For parishes, those that experienced more mobility were likely urban, had an institutionalized market prior to the Industrial Revolution, were less agrarian, and had lower income levels.
- Initial social structure was key, with mobility associated with having more gentry and fewer **yeomen**, which were a class of wealthier landowning peasants.
- Political characteristics are also significant: parishes including a member of Parliament are more mobile.

The implications of this research extend beyond historical curiosity to inform our understanding of social mobility more broadly. By focusing on the general population, the authors show that the Industrial Revolution fundamentally altered the relationship between family background and economic outcomes. Likewise, this work challenges narratives that focus solely on the Industrial Revolution's negative social impacts, suggesting instead that this period may represent a "dawn of liberty" for those previously entrenched in rigid social hierarchies. In doing so, this research contributes to ongoing debates about the relationship between economic development and social mobility, showing how a major economic transformation can weaken traditional status-based systems and create new pathways for advancement.

parish: a fundamental administrative unit for both religious and secular purposes, centered around a church and its priest, that served as the primary unit for managing local taxes, poor relief, and maintaining public infrastructure like roads and bridges

yeoman: a member of a rural middle class, ranking below the gentry, who typically owned and cultivated their own land, but also included skilled employees such as manor bailiffs, constables, and household servants

READ THE WORKING PAPER

NO. 2025-108 · AUGUST 2025

The Breakdown of the English Society of Orders: The Role of the Industrial Revolution

bfi.uchicago.edu/working-papers/the-breakdown-of-the-english-society-of-orders-the-role-of-the-industrial-revolution

ABOUT OUR SCHOLAR

James Robinson
Professor, Harris School of Public Policy and Department of Political Science; Fellow, Institute of African Studies at the University of Nigeria at Nsukka

THE UNIVERSITY OF CHICAGO
HARRIS SCHOOL OF PUBLIC POLICY

Written by David Fettig • Designed by Maia Rabenold

RESEARCH BRIEF • OCTOBER 2025

The Effects of Parental Income and Family Structure on Intergenerational Mobility: A Trajectories-Based Approach

Based on BFI Working Paper No. 2025-115, "The Effects of Parental Income and Family Structure on Intergenerational Mobility: A Trajectories-Based Approach," by Yoosoon Chang, Indiana University; Steven N. Durlauf, University of Chicago; Bo Hu, Indiana University; and Joon Park, Indiana University

Parental income and family structure during childhood and adolescence affect adult income, with these familial influences strongest in middle childhood and adolescence. The effects of income and family structure trajectories exhibit a complementary relationship during key developmental periods.

A rich literature within economics focuses on **intergenerational mobility**, or, the relationship between parents' socioeconomic status and their children's outcomes later in life. While much research focuses on how parents' average income shapes that of their child, research also suggests that timing matters across childhood. Moreover, joint consideration of family income and family structure is less common, despite evidence that family structure impacts a child's stability, emotional well-being, and mental health in ways that may interact with financial resources. Motivated by this, this paper tracks how parental income and family structure trajectories throughout childhood affect adult income.

Using data from over 1,000 families in the Panel Study of Income Dynamics, the authors develop

intergenerational mobility: the extent to which children's economic outcomes differ from those of their parents, measuring whether advantage or disadvantage passes from one generation to the next

Figure 1 · Dynamics of Overall, Structural, and Exchange Educational Mobility

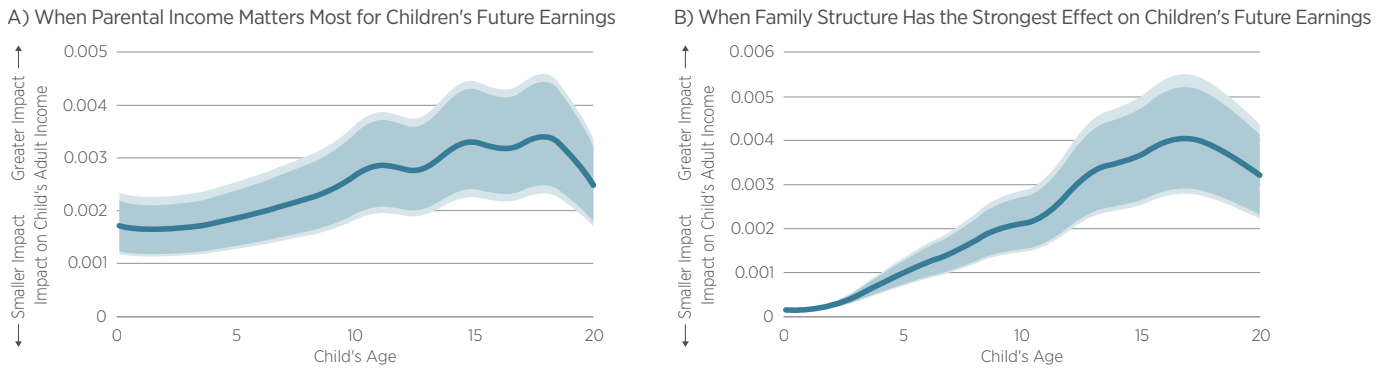

Note: This figure plots overall mobility (total probability children have different education than parents) with solid lines, exchange mobility (mobility after economic transitions stabilize) with dashed lines, and structural mobility (mobility due to temporary economic changes like education expansion) with vertical distance between lines, across generations. Steady-state mobility is represented by the rightmost points where overall equals exchange mobility. Shaded areas represent 95% confidence intervals.

a novel statistical framework that captures both the timing of family circumstances and their interactive effects across different developmental stages. Rather than examining income levels directly, the authors focus on income categories—classifying children into low-, middle-, and high-income brackets as adults. They define low income as below two-thirds of the median, high income as exceeding twice the median, and middle income as falling between these thresholds.

The authors analyze how parental income and family structure trajectories from birth to age 20 predict the probability that a child will end up in each of these three income classes, allowing them to identify when during childhood these family conditions have the strongest effects and whether they work as **complements** or **substitutes**. They find the following:

- While parental income benefits children at all ages, the effects are strongest during middle childhood and adolescence (ages 10-17) rather than early childhood. In other words, a child whose family has higher income during their teenage years is significantly more likely to achieve high adult income than one whose family had the same total income concentrated in early childhood.
- Family structure effects peak in adolescence. Children from two-parent households show modest advantages in early childhood, but these benefits become pronounced during middle childhood and peak around age 17. Living with both parents during the high school years provides the greatest boost to adult economic outcomes.
- Financial resources and family stability are complements, not substitutes. Rather than compensating for each other, parental income and two-parent family structure work synergistically. The benefits of higher family income are amplified when combined with family stability, and vice versa. This complementary relationship is strongest between parental income during middle to late childhood and family structure during late childhood.

These results underscore the importance of coordinated policies and interventions that address both economic and familial stability, particularly during critical developmental periods, to maximize the likelihood of positive outcomes for children. Importantly, the strength and timing of these effects may differ by gender, race and ethnicity, parental education, or other sociodemographic characteristics, suggesting an avenue for future research.

complements: two factors that work better together than separately, where the presence of one enhances the effectiveness of the other

substitutes: two factors that can replace each other, where having more of one reduces the need for the other

READ THE WORKING PAPER

NO. 2025-115 · AUGUST 2025

The Effects of Parental Income and Family Structure on Intergenerational Mobility: A Trajectories-Based Approach

bfi.uchicago.edu/working-papers/the-effects-of-parental-income-and-family-structure-on-intergenerational-mobility-a-trajectories-based-approach

ABOUT OUR SCHOLAR

Steven Durlauf
Frank P. Hixon Distinguished Service Professor; Director, Stone Center for Research on Wealth Inequality and Mobility, Harris School of Public Policy

THE UNIVERSITY OF CHICAGO
HARRIS SCHOOL
OF PUBLIC POLICY

Written by Abby Hiller • Designed by Maia Rabenold

RESEARCH BRIEF • OCTOBER 2025

The Impact of Language on Decision-Making: Auction Winners are Less Cursed in a Foreign Language

Based on BFI Working Paper No. 2025-124, "The Impact of Language on Decision-Making: Auction Winners are Less Cursed in a Foreign Language," by Fang Fu, Amazon Web Services; Leigh H. Grant, Exponent; Ali Hortaçsu, University of Chicago; Boaz Keysar, University of Chicago; Jidong Yang, Renmin University of China; and Karen J. Ye, Queen's University

Using a foreign language reduces the "winner's curse" in auctions, as bidders make more strategic decisions and are less likely to overbid when processing information in a non-native tongue. This effect largely disappears as bidders receive feedback across consecutive auctions, with both native and foreign language bidders converging to similarly poor decision-making patterns driven by observing others' overbidding behavior.

Research on the "foreign language effect" suggests that using a non-native language can change how individuals process information and make decisions. For example, people making moral judgments in a foreign language are more likely to approve utilitarian choices, such as pushing one person off a bridge to save five others, that they would reject when thinking in their native tongue. Similarly, people show reduced emotional responses when making judgments about risk and benefit in a foreign language, leading to more deliberative decision-making. These effects occur because foreign language use reduces reliance on quick, intuitive, emotionally driven judgments and promotes more deliberative thinking. In this paper, the authors test whether such language effects extend to financial decision-making in market settings.

The authors examine this question in the context of blind bidding auctions. They conduct a laboratory experiment in Beijing, China, with 357 native Mandarin Chinese speakers who know English as

Figure 1 · Difference Between Bids

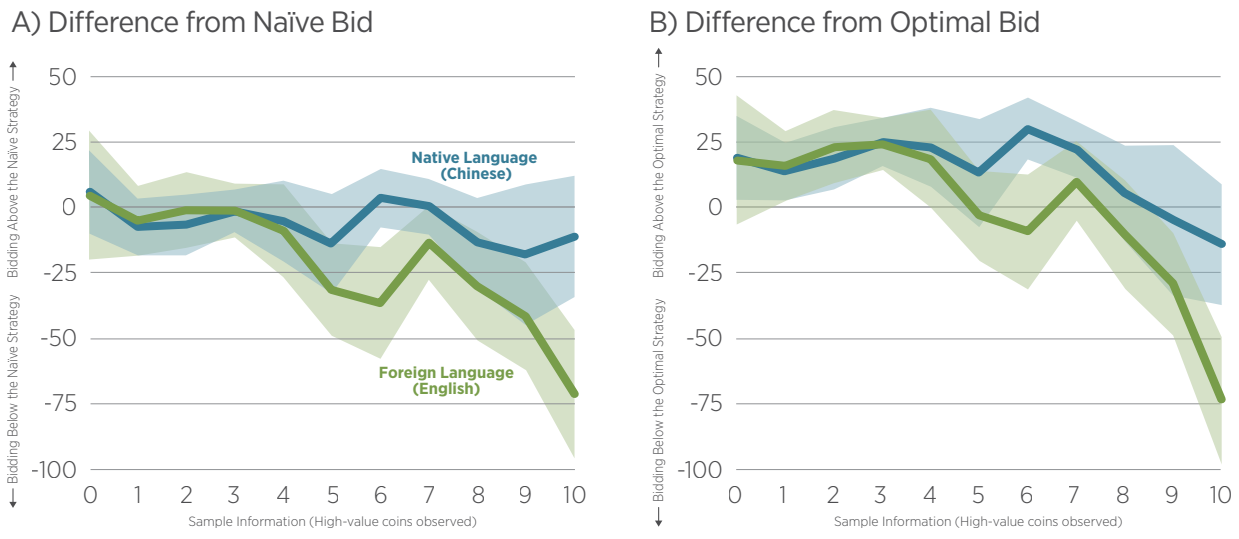

Note: These graphs show difference between individuals' bids and the Naïve (A) and optimal B) bid by language treatment in the first auction. Vertical bars are 95% confidence intervals.

a foreign language. Participants bid in ten rounds of auctions with the goal of winning jars filled with a random mix of high-value coins (1 RMB, worth about 15 cents) and low-value coins (0.1 RMB), while avoiding paying more than the value of the jar, or suffering the "winner's curse". Before each auction, each participant views a different sample of ten coins drawn from the jar to help estimate its total value. Participants are randomly assigned to complete the entire experiment in either Mandarin or English and in groups of 6 or 12 participants.

The authors measure how closely participants' bids match either a "naïve" strategy (bidding the expected value based on their sample of coins) or a theoretically optimal strategy that accounts for the risk of overbidding. They explore the role of language in influencing participants' bidding decisions, as well as the role of the number of bidders in an auction and response to feedback across auction rounds. The authors find the following:

- Bidders using a foreign language initially make more strategic decisions and are less susceptible to the winner's curse. In the first auction, 95% of native language bidders bid more than the value of the jar, compared to only 68% of foreign language bidders. Native language bidders bid significantly above the optimal level by an average of 16.05 RMB, while foreign language bidders did not deviate from the optimal bid.
- The foreign language effect disappeared across consecutive auctions as participants received feedback about others' bidding behavior. By the second through tenth auctions, bidders in both language treatments adopted similarly poor naïve bidding strategies, with winner's curse rates converging to around 85% in both groups.
- Feedback about previous winners' overbidding drove the convergence to poor decision-making. When auction winners overbid relative to the jar's actual value, other participants increased their own "cursedness" levels in subsequent auctions more than when winners did not overbid. This created a cycle where observing overbidding led to more overbidding.
- Decision making speed reflected the quality of choices across language treatments. Foreign language bidders took significantly longer to place bids in the first auction when they were making strategic decisions, but bidding speed converged between treatments as both groups adopted faster, more intuitive naïve strategies in later auctions.
- Group size effects were eliminated by foreign language use. Native language bidders in larger groups (12 participants) overbid significantly more than those in smaller groups (6 participants), replicating prior findings. However, foreign language bidders showed no difference in overbidding between large and small groups, suggesting the foreign language dampened competitive pressures.

These results have important implications for financial decision-making not only in auctions but more broadly in increasingly globalized market settings where participants frequently operate in non-native languages, suggesting both potential benefits and limitations of foreign language use in financial contexts. This interdisciplinary research was a joint effort of researchers from Economics and Psychology and was funded by NSF.

READ THE WORKING PAPER

NO. 2025-124 · SEPTEMBER 2025

The Impact of Language on Decision Making: Auction Winners are Less Cursed in a Foreign Language

bfi.uchicago.edu/working-papers/the-impact-of-language-on-decision-making-auction-winners-are-less-cursed-in-a-foreign-language

READ THE JOURNAL ARTICLE

JOURNAL OF ECONOMIC PSYCHOLOGY · OCTOBER 2025

The Impact of Language on Decision Making: Auction Winners are Less Cursed in a Foreign Language

sciencedirect.com/science/article/pii/S0167487025000571

ABOUT OUR SCHOLARS

Ali Hortaçsu
William B. Ogden Distinguished Service Professor of Economics, Kenneth C. Griffin Department of Economics

Boaz Keysar
William Benton Professor in Psychology, Department of Psychology

THE UNIVERSITY OF CHICAGO
KENNETH C. GRIFFIN DEPARTMENT OF ECONOMICS

THE UNIVERSITY OF CHICAGO
DEPARTMENT OF PSYCHOLOGY

Written by Abby Hiller • Designed by Maia Rabenold

RESEARCH BRIEF • OCTOBER 2025

Why Is Manufacturing Productivity Growth So Low?

Based on BFI Working Paper No. 2025-127, "Why Is Manufacturing Productivity Growth So Low?" by Enghin Atalay, Federal Reserve Bank of Philadelphia; Ali Hortaçsu, University of Chicago; Nicole Kimmel, Federal Reserve Bank of Philadelphia; and Chad Syverson, University of Chicago

Nearly all measured total factor manufacturing productivity growth since 1987, and its post-2000s decline, comes from a few computer-related industries. Conventional productivity growth statistics understate the manufacturing sector's productivity by failing to fully capture quality improvements. TFP growth is understated by 1.6 percentage points in durable manufacturing and 0.5 percentage points in nondurable manufacturing.

Official productivity statistics paint a picture of a manufacturing industry in decline. From 1987 to 2009, the Bureau of Labor Statistics (BLS) **total factor productivity** (TFP) index for manufacturing grew by 1.2% annually, outpacing the 0.9% growth rate of the overall private economy. But between 2009 and 2023, the pattern reversed. Manufacturing TFP fell slightly while private economy TFP continued to rise at 0.8% per year. This reversal is especially concerning given manufacturing's traditional role as a driver of innovation across the broader economy.

In this paper, the authors examine the recent slow growth in manufacturing productivity. They begin by documenting that nearly all measured TFP growth since 1987, and its post-2000s

Figure 1 · Productivity Growth: Manufacturing vs. Rest of Economy

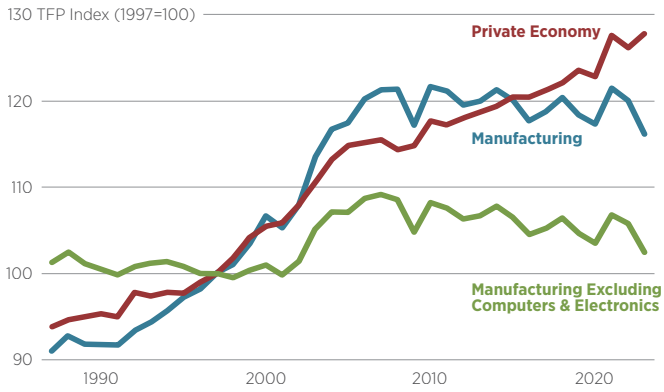

Note: This figure plots the cumulative contribution to manufacturing productivity growth for all industries other than Computer and Electronic Products Manufacturing.

total factor productivity: a measure of how efficiently a firm converts inputs (labor, capital, and materials) into output, capturing productivity improvements beyond changes in input usage

decline, comes from a few computer-related industries. While productivity growth slowdowns are observed in multiple manufacturing industries, most of the measured sector-wide stagnation is quantitatively explained by productivity changes in Computer and Electronic Products Manufacturing.

Building on this result, the authors next scrutinize whether official productivity measures accurately capture productivity for this highly innovative industry. Traditional TFP measures calculate productivity by comparing the **real** (inflation-adjusted) value of what industries produce to the real value of the inputs they consume. To make this calculation, statisticians must convert **nominal** dollar figures into real quantities, relying on **price indices** to strip out the effects of inflation.

These price indices must capture the frequent technological improvements within Computer and Electronic Products, or else risk undercounting the sector's production. For example, consider a smartphone that costs $800 today and $800 five years ago. Today's version likely has a better camera, faster processor, and more storage. When price indices fail to account for these quality improvements, they effectively treat a superior product as having the same price, missing the fact that the real (quality-adjusted) price has fallen dramatically.

This measurement error can have cascading effects, threatening to overstate inflation, understate real output growth, and make productivity gains appear smaller than they actually are. The authors hypothesize that this quality-adjustment problem is precisely what is happening in manufacturing statistics, causing conventional measures to understate productivity growth in highly innovative industries.

To test this hypothesis, the authors exploit a key institutional difference: the BLS invests more in quality adjustment for consumer price indices than producer price indices or import price indices. This is important because the last two price indexes are used to construct sector output and productivity, while consumer price indices are not. The authors compare consumer-facing prices (Personal Consumption Expenditures price index) against producer-facing prices (BEA gross output deflators and BLS import price indices) across 212 consumption categories and 414 commodities, using instances where consumer indices show steeper price declines than producer indices to identify unmeasured quality improvements.

To translate these price gaps into productivity estimates, the authors use **input-output tables**, which track what each industry buys from other industries, allowing them to trace the impacts of mismeasurement across the entire economy.

They find the following:

- Producer and import price indices understate quality growth. Comparing consumer-facing and producer-facing price indices for the same products reveals systematic gaps in rapidly innovating industries. For computers and electronics, consumer price indices show far steeper price declines than producer or import price indices, indicating that official producer deflators fail to capture quality improvements. This means measured inflation is overstated and real output growth is understated in these industries.
- Manufacturing productivity is systematically understated. Using an input-output framework to account for price mismeasurement in both outputs and inputs, the authors estimate that manufacturing TFP growth is understated by approximately 0.8 percentage points annually. The mismeasurement is concentrated in durable goods manufacturing (understated by 1.6 percentage points) with smaller effects

real: the value of money or interest rates after removing the effects of inflation, showing the true purchasing power

nominal: the stated or face value of something, like money or interest rates, without adjusting for inflation

price indices: a statistical measure that tracks how the average price of a basket of goods changes over time

input-output table: an economic accounting framework that maps which industries supply inputs to which other industries across the entire economy

in nondurable manufacturing (0.5 percentage points). No significant mismeasurement appears outside the manufacturing sector.

The manufacturing sector has undergone profound transformation over the past quarter century. Employment has collapsed by more than one-quarter since 1997 even as private nonfarm employment grew by more than one-quarter. The sector has become increasingly import-reliant—first from China, then from Vietnam and Mexico—more capital-intensive, and more robot-intensive. Assessments of manufacturing's evolution, and evaluations of the billions spent on federal support programs like SEMATECH, Manufacturing USA, and the CHIPS Act, hinge on properly measuring real output and productivity.

READ THE WORKING PAPER

NO. 2025-127 · SEPTEMBER 2025

Why Is Manufacturing Productivity Growth So Low?

bfi.uchicago.edu/working-papers/why-is-manufacturing-productivity-growth-so-low

ABOUT OUR SCHOLARS

Ali Hortaçsu
William B. Ogden Distinguished Service Professor of Economics, Kenneth C. Griffin Department of Economics

Chad Syverson
George C. Tiao Distinguished Service Professor of Economics, Chicago Booth

THE UNIVERSITY OF CHICAGO KENNETH C. GRIFFIN DEPARTMENT OF ECONOMICS

CHICAGO BOOTH THE UNIVERSITY OF CHICAGO BOOTH SCHOOL OF BUSINESS

Written by Abby Hiller • Designed by Maia Rabenold

ECONOMIC FINDING • JANUARY 2022

Inclusive Monetary Policy: How Tight Labor Markets Facilitate Broad-Based Employment Growth

Based on BFI Working Paper 2022-03, "Inclusive Monetary Policy: How Tight Labor Markets Facilitate Broad-Based Employment Growth," by Nittai K. Bergman, Tel Aviv University; David Matsa, Northwestern University; and Michael Weber, Chicago Booth

Expansionary monetary policy benefits the employment of workers with weak labor force attachment—Blacks, the least educated, and women—more in tight labor markets than in slack ones; this effect is economically sizeable and long lasting.

The Federal Reserve has recently emphasized the importance of understanding the labor market experiences of various communities when assessing its goal of maximum employment. Aggregate employment numbers, in other words, hide a lot of heterogeneity among groups, and the Fed has committed to addressing those differences.

However, there is little understanding of monetary policy's effects on different segments of the labor market. Does monetary policy, often described as a blunt instrument, impact different communities in different ways? If so, are there certain economic conditions under which the Fed can effectively target labor outcomes across different types of workers and demographic groups?

To address these and related questions, the authors of "Inclusive Monetary Policy: How Tight Labor Markets Facilitate Broad-Based Employment Growth," employed data from 895 local labor markets in the US between 1990 and 2019 to explore monetary policy's heterogeneous effects with respect to workers' race, education, and sex. Their key finding is that for demographic groups with low average labor market attachment—

Figure 1 • Impact of Monetary Policy on Employment Growth Over a One-Year Horizon
One-year horizon starts at different time periods following monetary policy rate change

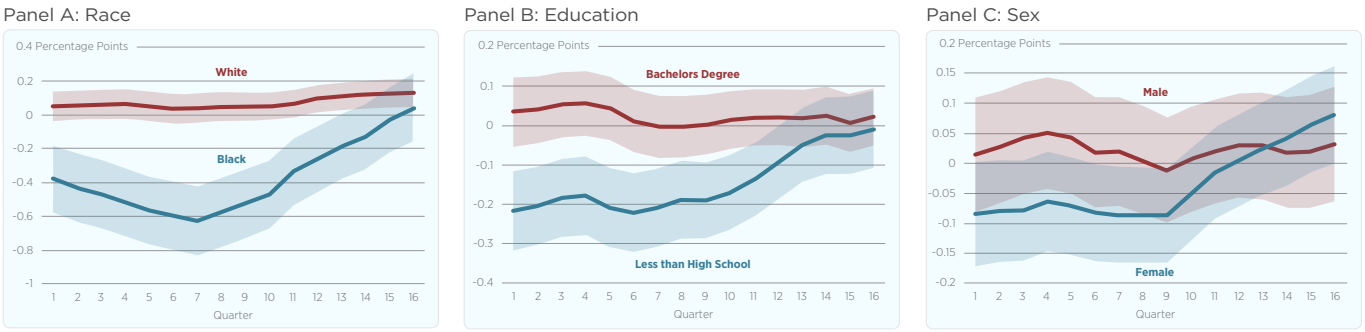

Note: This figure depicts the temporal dynamics of the differential impact of monetary policy on employment growth in tight versus slack labor markets. The figure shows the impact of monetary policy over a one-year horizon starting in different quarters following the monetary policy rate change for different demographic groups within three categories: race (Panel A), education (Panel B), and sex (Panel C). For each quarter, beginning one quarter to 16 quarters out, the figure plots the coefficient on the interaction term between the federal funds rate and the local prime age employment-to-population ratio in equation (4). Shaded areas present one standard deviation confidence intervals.

Blacks, the least educated, and women—monetary expansions have a larger effect on employment growth in tight labor markets. Importantly, this effect is economically large and persistent. For example:

- A one standard deviation drop in the federal funds rate in tight labor markets increases subsequent two-year Black employment growth by 0.91 percentage points, women's employment by 0.39 percentage points, and 0.37 percentage points for workers who did not complete high school.
- This additional impact of monetary policy in tight labor markets is sizable, corresponding to 9% and 18% of the mean employment growth rates for Blacks and high school non-completers over the sample period, respectively.
- Monetary policy's incremental effects on less-attached workers' employment growth in tight labor markets holds over time, peaking 7 to 9 quarters after interest rates decrease. (See Figure.)
- Finally, these effects are muted or non-existent for groups with stronger labor market attachment. For example, the point estimate for White employment growth is less than one quarter of the estimate for Blacks and not statistically significant.

This work suggests that sustained expansionary monetary policy, which tightens labor markets, facilitates robust employment growth among less-attached workers. Further, the Federal Reserve's recent change in its conduct of monetary policy from strict to average inflation targeting should benefit the employment of female, minority, and low skilled workers. At the same time, policy

READ THE WORKING PAPER

NO. 2022-03 · JANUARY 2022

Inclusive Monetary Policy: How Tight Labor Markets Facilitate Broad-Based Employment Growth

bfi.uchicago.edu/working-paper/202203

ABOUT OUR SCHOLAR

Michael Weber
Associate Professor and Fama Faculty Fellow, Chicago Booth

CHICAGO BOOTH
THE UNIVERSITY OF CHICAGO BOOTH SCHOOL OF BUSINESS

THE UNIVERSITY OF CHICAGO
BECKER FRIEDMAN INSTITUTE
FOR ECONOMICS

KEEP UP WITH BFI

bfi.uchicago.edu