IMF RESEARCH perspectives

Volume 27

December 2025

CONTENTS

IMF RESEARCH
perspectives

GUEST EDITORS
Sneha Agrawal and Flavien Moreau

EDITOR-IN-CHIEF
Petia Topalova

ASSISTANT EDITOR
Patricia Loo

IMF CSF CREATIVE SOLUTIONS DIVISION
Art Direction/Layout: Laura Spofford
Cover Illustration: Stefan Lipsky
Project Management: Razan Elbaba

CONTRIBUTORS
Jaedo Choi
George Cui
Adriano Fernandes
Kailin Gao
Nils Haakon Lehr
Nan Li
Roldolfo Rigato
Younghun Shim
Liang Zhang
Robert Zymek

IMF Research Perspectives—the IMF online bulletin with news on research—is an annual publication in English and is available exclusively online free of charge.

Editorial correspondence may be addressed to:

INTERNATIONAL MONETARY FUND
Attn: Editor, IMF Research Perspectives
700 19th Street, NW
Room HQ1 9-253
Washington, DC 20431
USA
resbulletin@IMF.org

ISBN=9798229030847
DOI=10.5089/9798229030847.053
ISSN=2708-3535

NOTE FROM THE GUEST EDITORS

Average global growth in the next five years of the 2020s is forecast to be slowest of any decade since the 1960s. Should this forecast materialize, the repercussions could be wide-ranging: Advanced economies may not see the standards of living of their citizens rise in a meaningful way, those in lower-income countries may fall further behind, and tackling debt could become even more challenging. Against this backdrop, this issue of *IMF Research Perspectives* highlights five novel studies by IMF staff that examine the factors that may be contributing to the slowdown in potential growth and point to "new levers of growth."

The articles featured in this issue study the accumulation of factors of production, such as labor and capital; the efficiency with which these are used; and innovation, which drives technological progress and ultimately underpins long-term growth. The evidence presented in this issue draws from the experiences of countries from a wide spectrum of the income distribution—including China, India, Korea, and the United States—and emphasizes how well-designed policies and emerging technologies can overcome frictions faced by firms and workers to unleash economic development and growth. This issue does not touch upon AI, a potentially powerful new lever for growth. We refer our readers to the May 2024 issue of *IMF Research Perspectives*, "New Technologies, Digitalization, and AI," which was dedicated to this topic.

We are honored to present an in-depth interview with Chris Papageorgiou, mission chief for Bangladesh, former Division Chief in the IMF's Research Department, and *IMF Research Perspectives'* longest-serving editor, covering his research on growth and development; and his policy experience as mission chief.

~ Sneha Agrawal and Flavien Moreau

SIGN UP for new issues of *IMF Research Perspectives* and a variety of other IMF publications; sign up for electronic notification at IMF.org/external/cntpst. Past issues are available at IMF.org/researchbulletin.

AN INTERVIEW WITH

CHRIS PAPAGEORGIOU

ASSISTANT DIRECTOR,
ASIA AND PACIFIC DEPARTMENT, IMF

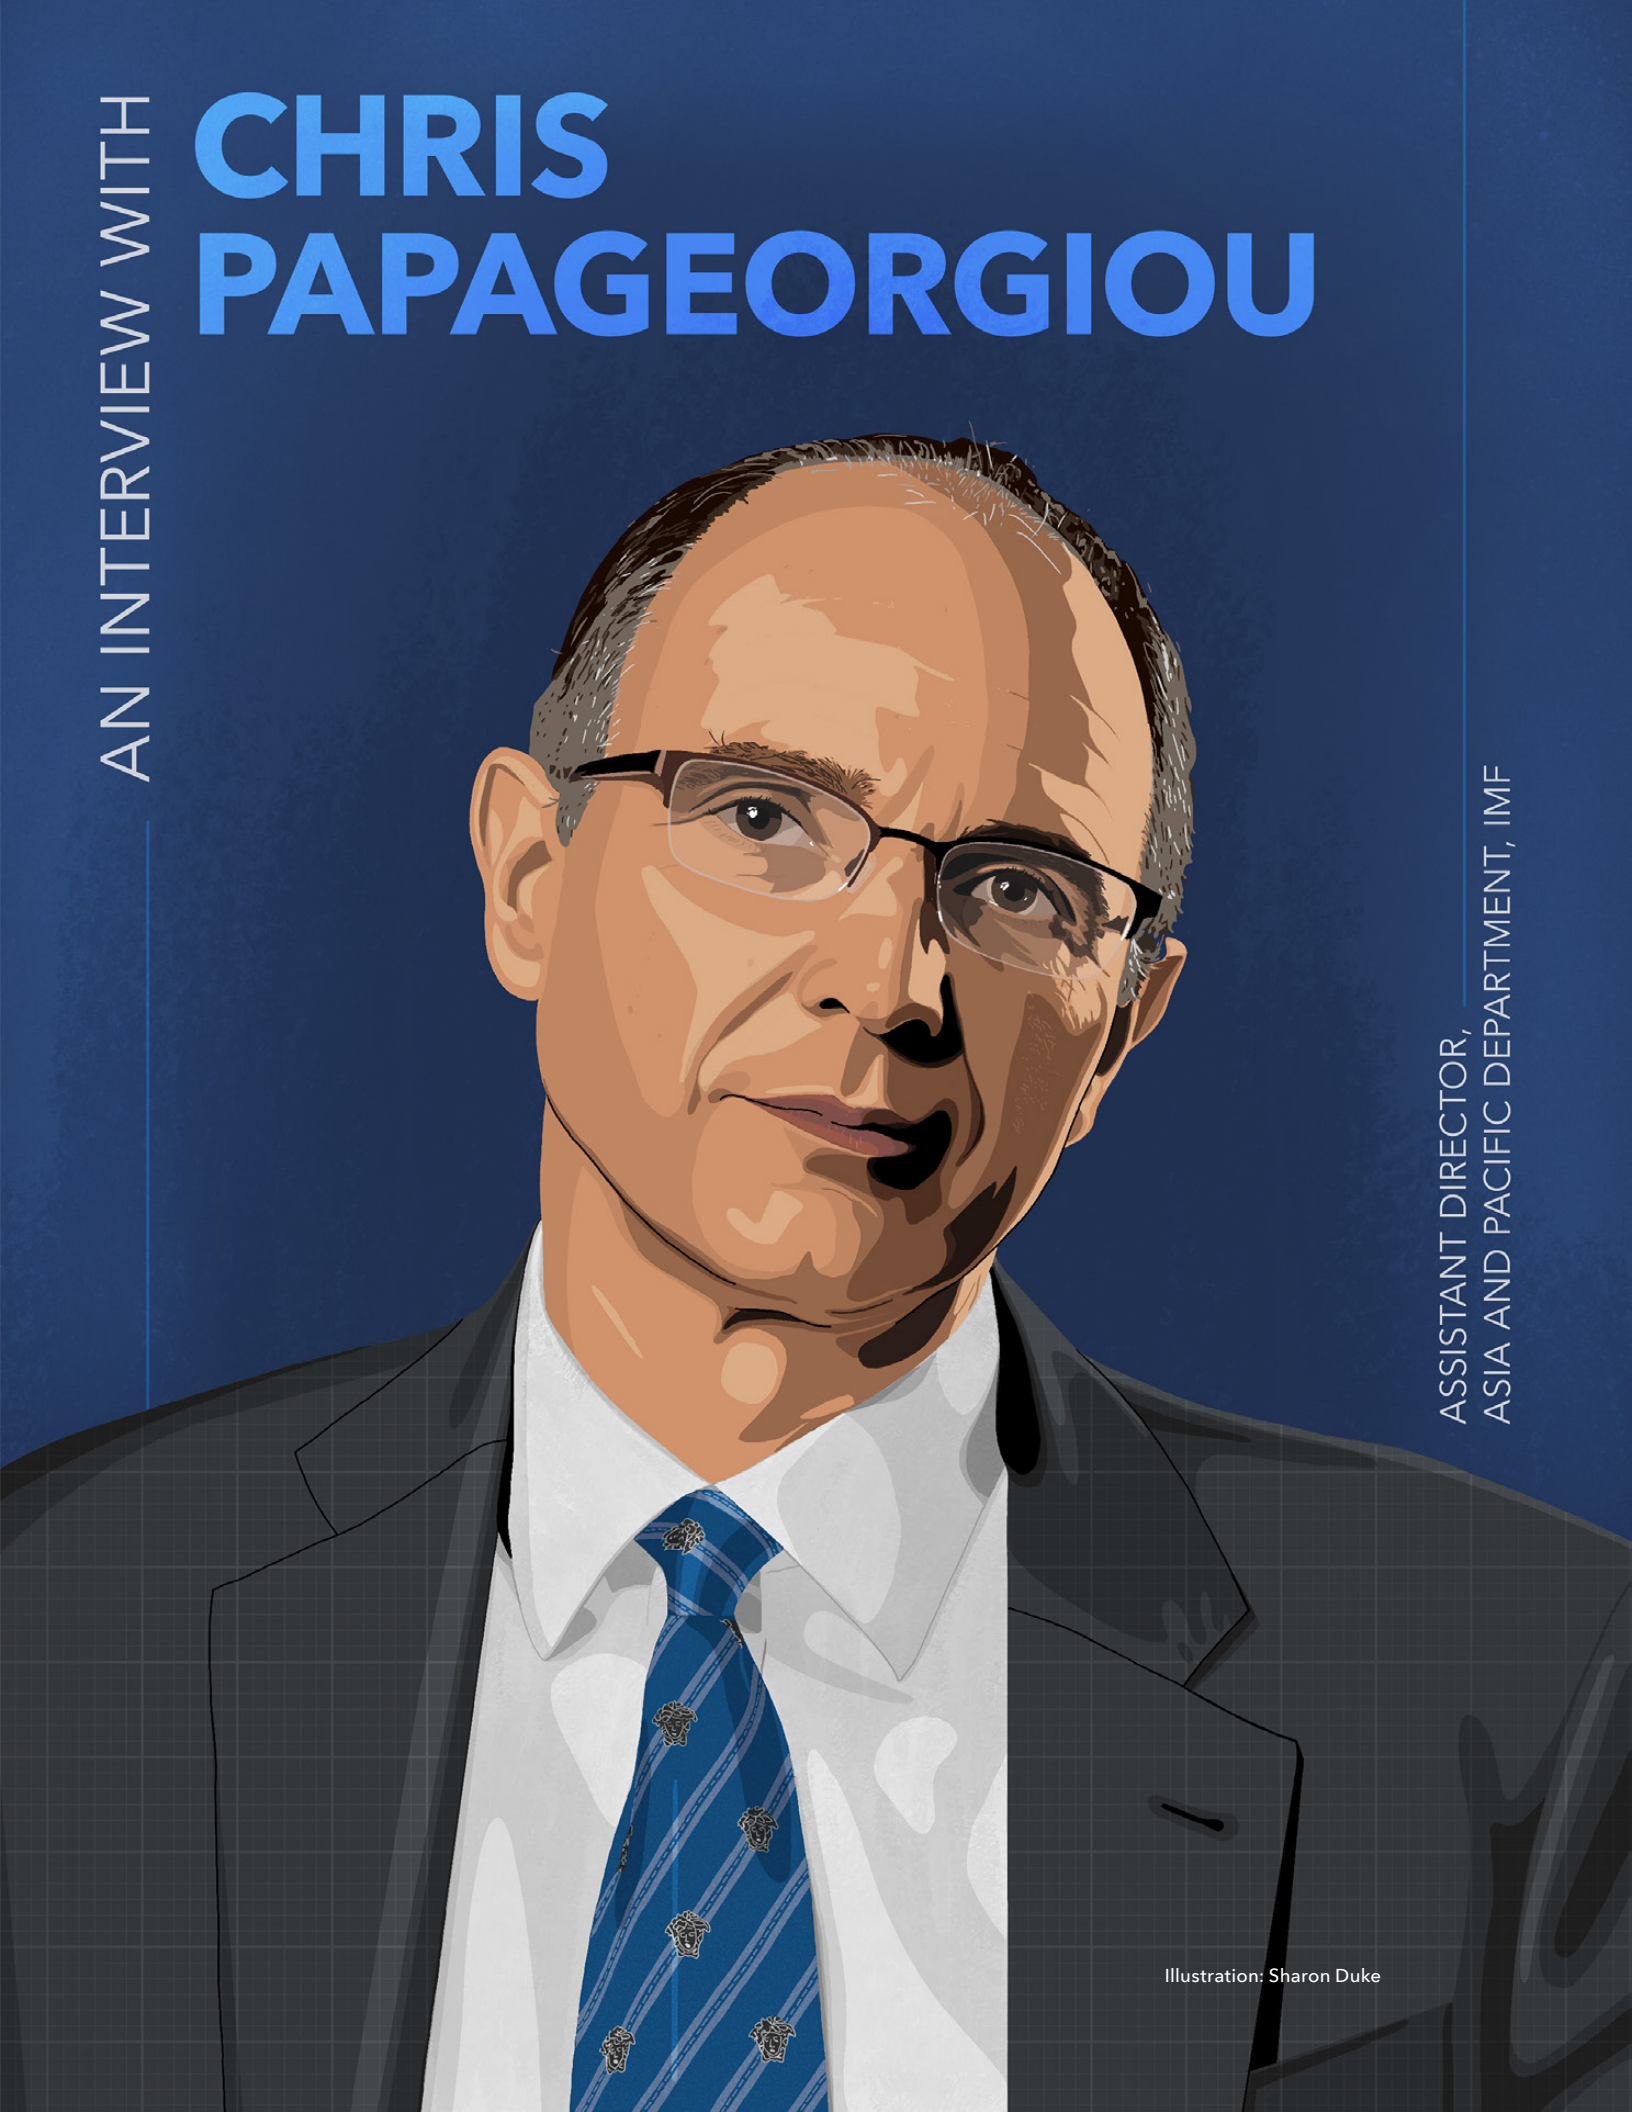

Illustration: Sharon Duke

This year's interview features Chris Papageorgiou, Assistant Director of the IMF's Asia Pacific Department and former editor of *IMF Research Perspectives*. The conversation focuses on the search for new growth levers in a rapidly changing global economy. Chris shares valuable insights from his experience as mission chief for Bangladesh, addressing the complexities of development amid geopolitical shifts, disruption of traditional value chains, and the transformative impact of climate change and AI. The interview delves into policy responses for inclusive and sustainable growth in emerging market and low-income economies.

Here is a brief excerpt from the interview:

Sneha Agrawal: What are your views on the most promising "*new* levers of growth," especially for low-income and emerging market economies in today's context?

Chris Papageorgiou: Over the past six decades, economic thought has equipped us with a rich arsenal of growth levers—from investment and education to innovation and the strengthening of institutions. This comprehensive toolkit has been invaluable in shaping effective policy advice. Yet today we stand at the threshold of a new era defined by revolutionary forces such as artificial intelligence and the urgent realities of climate change—factors poised to redefine global growth for generations to come.

However, as we engage with authorities on the ground, a striking new challenge emerges: The world is shifting away from the era of globalization that has long shaped our economic thinking and policy advice. Instead, we're entering a new phase where success depends on being plugged into the right networks, forming regional ties, and thinking strategically about where you belong in a more fractured economic landscape. It's no longer enough to just do the "right" things and expect markets to reward you. The old rules are giving way to a tougher, more complex reality—meaning policymakers have to be sharper and more strategic than ever before.

Watch the interview here

WHEN CAPITAL CONSTRUCTION SPEEDS UP OR SLOWS DOWN

New Evidence on Endogenous Time to Build

Adriano Fernandes
afernandes@IMF.org

Rodolfo Rigato
rodolfo.dinis_rigato@ecb.europa.eu

Illustration: Stefan Lipsky

How long does it take to build a factory, a power plant, or a bridge? The conventional wisdom in economics treats construction time as a fixed technological constraint—a steel plant takes three years to build, period. But what if firms could choose to speed up or slow down construction based on economic conditions?

Our new research using detailed project-level data from India shows that time to build is indeed a choice variable, with profound implications for how we think about investment dynamics and economic policy (Fernandes and Rigato 2025).

Construction Times in Economics

Physical capital doesn't materialize overnight. The Pentagon's construction during World War II illustrates this reality. Originally budgeted at $35 million with a standard timeline, urgent wartime needs led to round-the-clock construction that nearly doubled costs to $63 million but dramatically shortened completion time. More recently, Operation Warp Speed demonstrated how vaccine development and construction of manufacturing facilities could be accelerated when circumstances demanded it.

Despite these anecdotes, economic models have traditionally assumed that construction schedules are immutable—if a project takes three years, that's simply how long it takes. This assumption matters because investment is a key driver of economic fluctuations, and construction delays create a wedge between when firms decide to invest and when productive capacity comes online.

Measuring Time to Build at Scale

To move beyond anecdotes, we leveraged a unique dataset tracking over 100,000 large capital projects in India from the 1990s through 2022. These projects span everything from manufacturing plants to airports, from irrigation canals to office buildings. India provides an ideal laboratory as the world's third-largest construction market, with $1.7 trillion in anticipated investment by 2030.

The data reveal striking patterns (Figure 1). The median project takes 2.5 years to complete, but the distribution has a long right tail—10 percent of projects take over a decade. This heterogeneity isn't random: Transportation and infrastructure projects take longest, while retail and IT facilities are completed fastest. Larger projects take proportionally longer, with a 1 percent increase in project size associated with a 0.18 percent increase in construction time. Importantly, we can explain only about 30 percent of the variation in construction times using industry, location, and other observable project characteristics.

Figure 1. Distribution of Construction

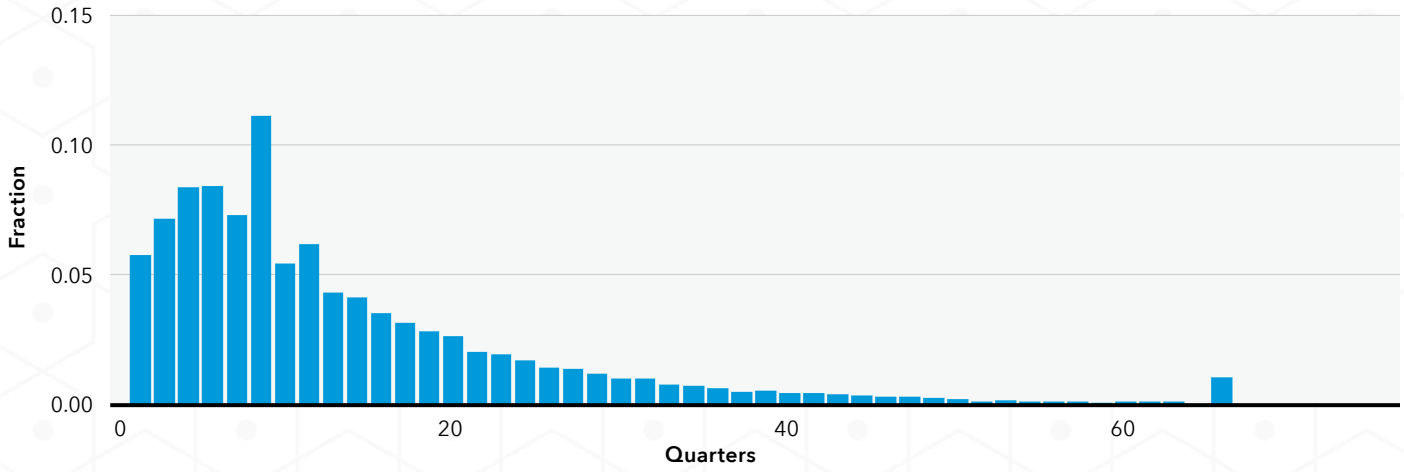

Sources: CapexDx March 2022 Vintage and Authors' Calculations.
Note: The figure plots the histogram of gestation lags for all completed projects in the March 2022 Capex Dx vintage. We measure a gestation lag from the first date that a project appears in the data (its announcement date) to its eventual completion. We winsorize the distribution at the 1st and 99th percentiles for visualization.

Speeding Up

To test whether firms can adjust construction speed, we need variations in economic conditions that affect firms (such as, for example, their ability to obtain credit) but are unrelated to construction's physical requirements. India's Asset Quality Review provides exactly such a shock.

Starting in 2015, the Reserve Bank of India conducted comprehensive audits of commercial banks, forcing them to recognize bad loans and set aside additional provisions. Banks more exposed to these audits cut lending sharply. Importantly, the audits were regulatory interventions during a non-crisis period, making them plausibly unrelated to construction fundamentals. We trace this banking shock to individual projects by linking firms to their lenders. We show that firms borrowing from more affected banks faced a credit crunch.

Economic intuition suggests that credit-constrained firms should cut investment across the board. We find something more nuanced: Firms accelerate ongoing projects while starting fewer new ones. Specifically, a one-standard-deviation increase in exposure to the credit shock increases the quarterly probability of completing ongoing projects by 6 percent, but reduces new project starts by 7.5 percent.

This pattern makes economic sense once we recognize that firms face a fundamental trade-off. Speeding up construction is costly—it requires overtime pay and expedited deliveries, and carries coordination challenges. But completing projects faster also brings benefits: The firm starts earning returns sooner. When credit tightens and firms' discount rates rise, the benefit of early completion can outweigh the acceleration costs, especially for projects already near completion.

Not All Projects Are Equal: The Role of Maturity

Our theoretical framework predicts that acceleration should be strongest for projects closest to completion (Figure 2). To test this, we examine whether the credit shock affects projects differently based on their stage of construction.

The data strongly support this prediction. Projects estimated to be in later stages of construction show significantly stronger acceleration in response to the credit shock. This heterogeneity provides compelling evidence that our results reflect active management of construction schedules rather than some unobserved shock to construction productivity.

Implications for Economic Dynamics

These findings change how we should think about investment dynamics. When time to build is endogenous:

(1) Investment becomes more volatile: Firms can partially front-load or back-load construction spending in response to shocks. Our model simulations suggest that endogenous time to

Figure 2. Heterogeneous Effects by Project Stage

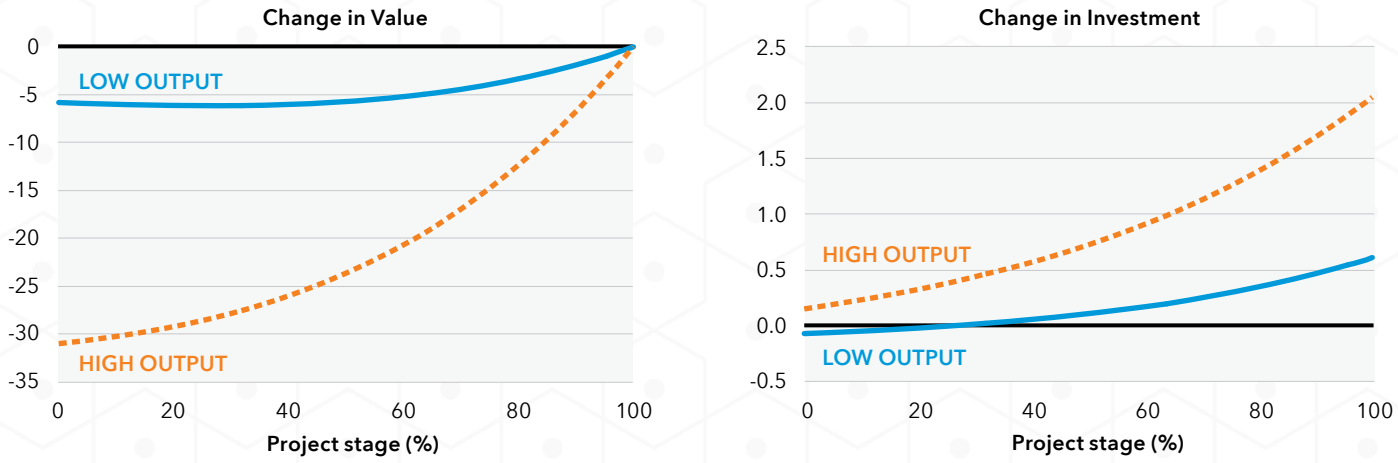

Source: Authors' Model Computations.
Note: The figure plots the response of project value and investment policy following a shock to discount rates in percent deviations from steady state values. The dashed orange (solid blue) line represents projects that yield large (small) output once completed.

build can amplify investment fluctuations by about 30 percent compared with models with fixed construction times.

(2) **Responses depend on the project pipeline:** The aggregate investment effect of any shock depends on the distribution of projects across construction stages. Aggregate investment in an economy with many near-complete projects will respond differently from one dominated by newly started projects.

(3) **Accounting versus productive investment diverge:** Traditional investment measures capture spending, but what matters for productivity is completed projects. These can move in opposite directions when construction speeds adjust.

Rethinking Monetary Policy

The state-dependent nature of investment responses has crucial implications for monetary policy. When central banks raise interest rates, the investment response depends critically on the maturity distribution of ongoing projects (Figure 3).

Consider two scenarios:

- **Many mature projects:** An increase in interest rates would accelerate near-complete projects (as firms rush to finish and start earning returns) while dramatically reducing new starts. Investment might temporarily rise before collapsing.
- **Many young projects:** An increase in interest rates would slow investment, producing a more traditional contractionary effect.

This state-dependence means that monetary policy effectiveness varies over the construction cycle. Central banks may find their tools less potent when many projects are near completion—precisely when inflationary pressures from capacity constraints might be strongest.

Fiscal Policy as a Timing Tool

Investment tax credits and subsidies take on new meaning when construction timing is flexible. These policies don't just change how much firms invest—they can change when investment occurs.

Our framework shows that temporary investment tax credits create powerful incentives to accelerate ongoing projects. Firms rush to complete projects before subsidies expire, front-loading investment spending (Figure 4). This provides fiscal authorities with a precise tool to shift aggregate demand across time.

For structural policies like green energy transitions, acceleration incentives involve a clear trade-off: faster deployment of new technologies versus higher costs. Policymakers must weigh whether getting renewable energy infrastructure online six months sooner justifies potentially 20–30 percent higher construction costs.

Figure 3. State-Dependent Monetary Policy

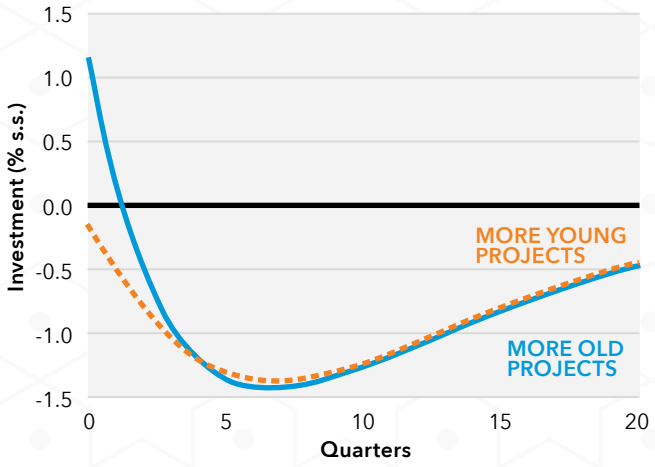

Source: Authors' Model Computations.

Note: The figure illustrates the response of aggregate investment over time to a contractionary interest rate shock, holding the price of capital fixed. The dashed orange (solid blue) line shows the response when the initial maturity distribution skews towards younger (older) projects.

Figure 4: Tax Policy Reshuffles Investment Dynamics to Within-Project Reallocation of Resources over Time

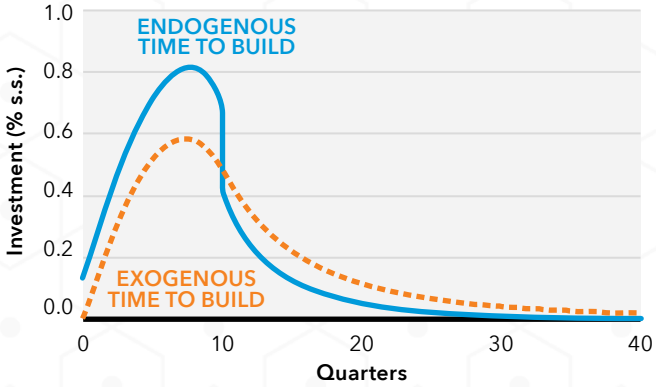

Source: Authors' Model Computations.

Note: The figure illustrates the response of aggregate investment over time to an unanticipated introduction of an investment tax credit at t=0 lasting for ten quarters, allowing the price of capital to vary. The dashed orange (solid blue) line shows the response under exogenous (endogenous) time to build.

Practical Implications for Investment Planning

For corporate managers and investors, our findings suggest several concrete implications:

(1) Capital budgeting must consider timing flexibility: Traditional net present value (NPV) calculations assume fixed project schedules. But the option to accelerate or delay completion has substantial value, especially for long-duration projects.

(2) Project portfolios matter: Firms with diverse project portfolios—some near completion, others just starting—have more flexibility to respond to shocks. This construction-stage diversification provides a buffer against credit market disruptions.

(3) Beware of accounting illusions: Reported investment spending may rise even as firms are retrenching, simply because they're accelerating near-complete projects while canceling planned ones. Cash flow analysis must look beyond current spending to the pipeline of future commitments.

The Bigger Picture

The ability to adjust construction timing reveals a deeper truth about modern economies: What appear to be technological constraints often hide important economic choices. Just as just-in-time manufacturing revolutionized how we think about inventory management, recognizing time to build as endogenous changes how we understand investment dynamics.

This insight is particularly relevant for emerging markets undertaking massive infrastructure build-outs. Countries like India, Indonesia, and Vietnam are simultaneously building power plants, transportation networks, and manufacturing facilities. The ability to strategically accelerate or delay these projects provides valuable macroeconomic flexibility.

For advanced economies pivoting toward green infrastructure, endogenous construction timing offers both opportunities and challenges in the face of policy uncertainty. Acceleration can help meet climate targets sooner but at higher costs. Understanding these trade-offs is essential for designing effective industrial policy.

Looking Forward

Our findings open several important avenues for future research and policy development:

(1) Optimal contract design: How can construction contracts better balance flexibility with cost control? Current fixed-price contracts may inefficiently restrict timing adjustments.

(2) Cross-country differences: Do institutional factors like labor market flexibility or regulatory frameworks affect how much construction timing can adjust?

(3) Financial innovation: Can financial markets develop instruments that help firms manage the risks of construction timing decisions?

As countries worldwide grapple with infrastructure needs, green transitions, and postpandemic recoveries, the ability to strategically manage construction timing becomes ever more valuable. Policymakers who understand these dynamics can design more effective interventions. Firms that master these choices can better navigate economic volatility. Concrete and steel may be solid, but the timeline for putting them in place is remarkably fluid.

REFERENCE

Fernandes, Adriano, and Rodolfo Rigato. 2025. "K Wasn't Built in a Day." IMF Working Paper 25/078, International Monetary Fund, Washington, DC.

MISALLOCATION of TALENTS

How Gender-Based Occupational Barriers Constrain Aggregate Productivity

Nan Li
nli@IMF.org

Robert Zymek
rzymek@IMF.org

Liang Zhang
zhangpw@bc.edu

Illustration: Michelle Martin

Over recent decades, the labor market for women in advanced economies has undergone a remarkable transformation.

In 1965, US women aged 25–34 were six times less likely than men to be in formal employment. Among those who did work, nearly 40 percent were in a single occupation: clerical support roles, and even within that occupation, women earned just 66 cents for every dollar earned by men. By 2015, US women's labor force participation rate had tripled, their employment was more evenly distributed across occupations, and wage gaps had narrowed—with female clerical support workers earning 90 cents on the dollar earned by their male counterparts.

A growing body of literature suggests that this transformation was not incidental to US economic performance. An influential study by Hsieh and others (2019) argues that the decline in occupational barriers for women—and for Black workers—may account for 20 to 40 percent of per capita income growth in the US between 1960 and 2010. These gains came not only from higher labor force participation, but also from better *talent allocation*: As barriers fell, workers were able to pursue occupations better aligned with their skills and preferences, raising aggregate productivity.

Building on this insight, our research (Li, Zhang, and Zymek, forthcoming) extends the framework of Hsieh and others (2019) to a cross-country setting. Using 2015 harmonized income microdata for 43 economies, we estimate the magnitude of gender-based occupational barriers implied by observed labor market outcomes (incomes and participation across various occupations) and quantify their economic impact. Our measure captures frictions beyond conventional wage gaps—such as discriminatory practices, social norms, or costs to human capital accumulation—and allows consistent comparison across countries and time.

The results reveal stark disparities between high-income economies and emerging market and developing economies (EMDEs) (Figure 1). In Egypt and India, the data-implied barriers facing women are almost three times as high as in the US in 2015—levels comparable to the US in the 1960s, when gender-based occupational segregation was widespread. At the same time, several EMDEs in our sample exhibit some of the lowest estimated barriers. Many of these are former centrally planned economies where participation in the workforce regardless of gender was historically mandated.

Why do these barriers matter? Beyond discouraging participation, they result in a misallocation of talents preventing women from entering

Figure 1. Composite Occupational Barriers for Women across Economies in 2015

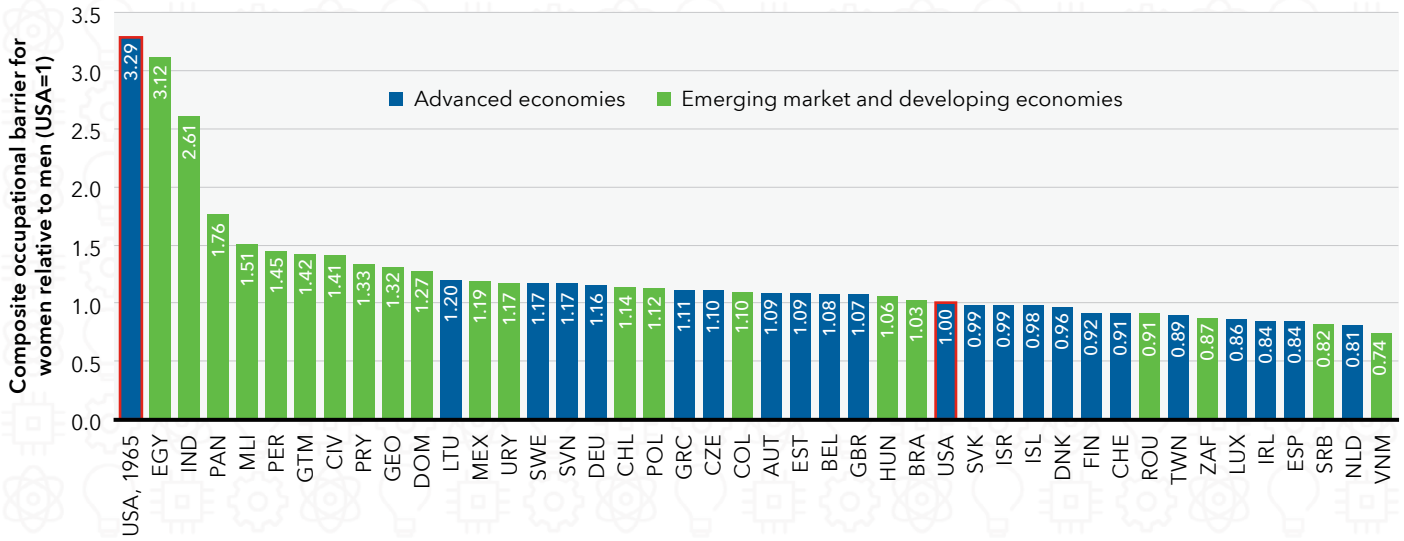

Sources: Li, Zhang, and Zymek, forthcoming; and Luxembourg Income Study database.
Note: Data labels in the figure use International Organization for Standardization (ISO) country codes.

occupations where their comparative advantage would be highest. Our counterfactual simulations show that lowering female occupational barriers to 2015 US levels would yield double-digit gains in per capita market incomes for some major EMDEs, with most of the gains driven by productivity improvements. This suggests that policies aimed at reducing occupational barriers for women could be a powerful source of growth and international income convergence.

Measuring the Barriers across Countries

Traditional measures of gender inequality in the labor market—such as wage gaps or workforce participation—tell only part of the story. On one hand, a country might have high female participation yet still harbor significant barriers if women are steered away from high-paying professions. On the other hand, a modest wage gap might mask severe occupational segregation, since only the most capable women may overcome high barriers to enter certain occupations. These women often earn high wages, which obscures the extent of exclusion faced by others: Consider a comparison of the earnings of Sandra Day O'Connor or Ruth Bader Ginsburg with their average male contemporary in the legal profession.

Our measure of labor market barriers facing women builds on a choice-theoretic model along the lines set out in Hsieh and others (2019). Individuals choose occupations and education investment early in life to maximize their lifetime utility, based both on group-specific and idiosyncratic talents and preferences. These choices may be distorted by group-specific barriers that act like hidden taxes, reducing the attractiveness of certain occupations for women relative to men. These barriers can be interpreted broadly—they may reflect discriminatory policies, informal institutions, or social norms that prevent individuals from pursuing the occupations where they could be most productive. By comparing observed patterns of occupational sorting and earnings between women and men, our approach infers the magnitude of these barriers while controlling for unobserved differences in talent and preferences.[1] The barrier estimates are comparable across

1 A country's barrier measure is a selection-adjusted gender wage gap. We first calculate the geometric mean ratio of women's to men's earnings, corrected for relative participation odds. These gaps are then aggregated across occupations based on earnings share.

Figure 2. Composite Occupational Barriers for Women and Gender Inequality Indicators

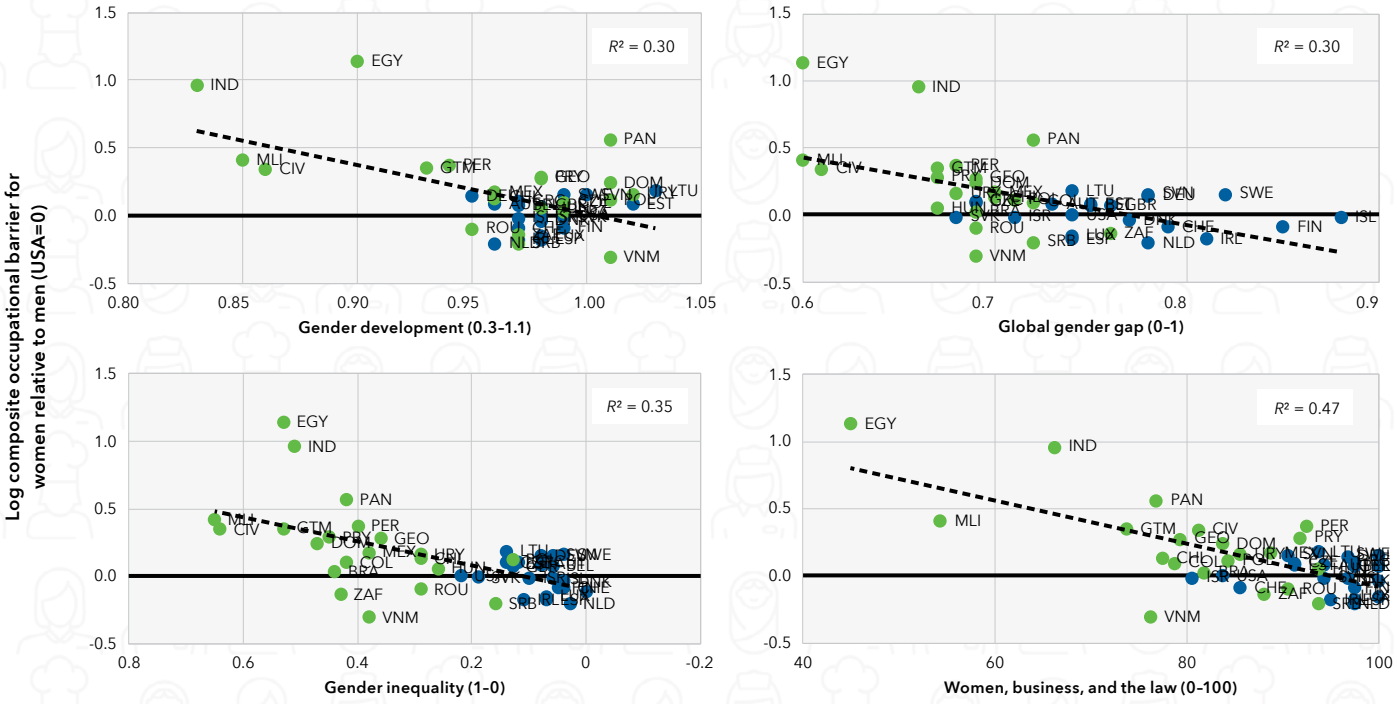

Sources: Li, Zhang, and Zymek, forthcoming; United Nations Development Programme 2015; World Bank 2016; and World Economic Forum 2015.
Note: Data labels in the figure use International Organization for Standardization (ISO) country codes.

countries under the crucial but plausible assumption that the innate comparative advantages of women relative to men—absent institutional and economic distortions—are the same everywhere.

Our empirical analysis relies on data from the Luxembourg Income Study, which provides the largest harmonized microdata on individual and household labor market outcomes. Drawing on income surveys from 52 economies across all regions, the study ensures comparability across countries through standardized occupation codes, income definitions, and employment measures.

The resulting measure of female occupational barriers is strongly correlated with several established indices of gender inequality (Figure 2). Our composite barriers align closely with both outcome-based measures—such as the Gender Development Index (GDI, UNDP 2015), Global Gender Gap Index (GGGI, WEF 2015), and Gender Inequality Index (GII, UNDP 2015)—as well as an institution-based metric such as the Women, Business and the Law Index (World Bank 2016). The *R*-squared values from regressing our measure of occupational barriers on these indices range from 0.30 to 0.47, underscoring that our model-based approach captures meaningful variation in gender disparities as reflected in established benchmarks.

Gains from Lowering Barriers: Improved Talent Allocation and Participation

Since our measure of barriers is derived from a formal economic model, it readily lends itself to counterfactual analysis. Using the US as a benchmark, we simulate the income gains countries could achieve if their female occupational barriers fell to US 2015 levels (Figure 3).

Almost all economies in our sample experience some income gains in this scenario, closely correlated with the size of their barriers relative to the US, identified in Figure 1. The largest gains accrue to countries with the highest barriers. A significant number of EMDEs in our study benefit from income gains about 10 percent or more. This includes the large emerging markets of Egypt and India, whose real per capita GDP increases by 26 percent and 24 percent, respectively. Gains for advanced economies (AEs) are relatively modest: The Czech Republic, Germany, Iceland, and Lithuania—among the highest-barrier AEs—could expect income increases of about 5 percent. Nonetheless, against the backdrop of declining fertility rates and aging populations, even these moderate gains could be a powerful lever to boost medium-term growth and sustain labor force participation in AEs.

Figure 3 also highlights the two channels through which the income gains are realized. First, there is a familiar participation channel (light-shaded bar segments): With lower barriers, more women leave the "home" sector for market work, which raises per capita income. Second, and more important, is a productivity effect (dark-shaded bar segments): With the leveling of occupational barriers, women can sort into the occupations that better reflect their true comparative advantage. For most economies in our sample, the productivity gains from improved talent allocation across market occupations make up the largest portion of the hypothetical GDP increase. For the median sample economy, about two-thirds of income gains are derived from improved productivity, and about one-third from increased participation.

Role of Sectoral Factors

Since our measure of female occupational barriers aggregates across occupations using earnings weights, it is useful to disentangle the two underlying drivers of high composite barriers in a given country. First, the country may exhibit high within-occupation barriers. Second, even if within-occupation barriers are moderate, the overall barrier may still be high if the country's economic structure is tilted toward occupations that are female-unfriendly. A shift-share decomposition reveals that for the large EMDEs with the highest barriers, both forces play a role. That is, these economies have higher barriers than the US at the occupation level, but their composition of economic activity is also unfavorable to female participation. The countries with the lowest overall barriers tend to achieve that status primarily through lower within-occupation barriers.

Figure 3. Baseline Real Per Capita Income Gains from Lowering Occupational Barriers

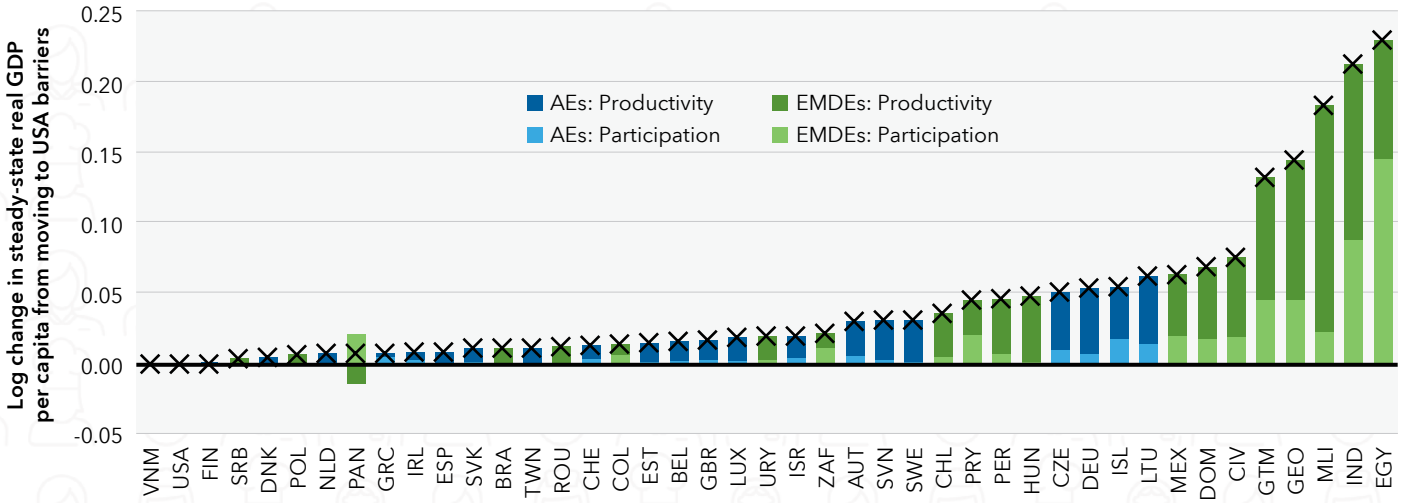

Sources: Li, Zhang, and Zymek, forthcoming; and Luxembourg Income Study database.
Note: AEs = advanced economies; EMDEs = emerging market and developing economies. Data labels in the figure use International Organization for Standardization (ISO) country codes.

A related concern is whether some high estimated barriers reflect genuine differences in comparative advantage, especially in sectors that require substantial physical strength. To address this, we explore an alternative assumption that "brawn-intensive" occupations—such as skilled agriculture, crafts, and machine operation—may exhibit large selection-adjusted wage gaps that reflect male comparative advantage rather than barriers. Under this alternative, the variation in barriers across economies narrows somewhat, but the ranking of economies by the size of their barriers remains virtually unchanged. Furthermore, eliminating all barriers under this alternative assumption still yields sizable per capita GDP gains (~9.3 percent).

Welfare Effects

Figure 4 presents the steady-state per capita welfare effects of reducing occupational barriers. Although overall gains (black dots) are broadly in line with the GDP effects, they tend to be slightly smaller. This reflects the trade-off between higher market consumption (reflected in higher GDP) and lower consumption of home services:

Figure 4. Baseline Welfare Gains from Lowering Occupational Barriers, by Gender

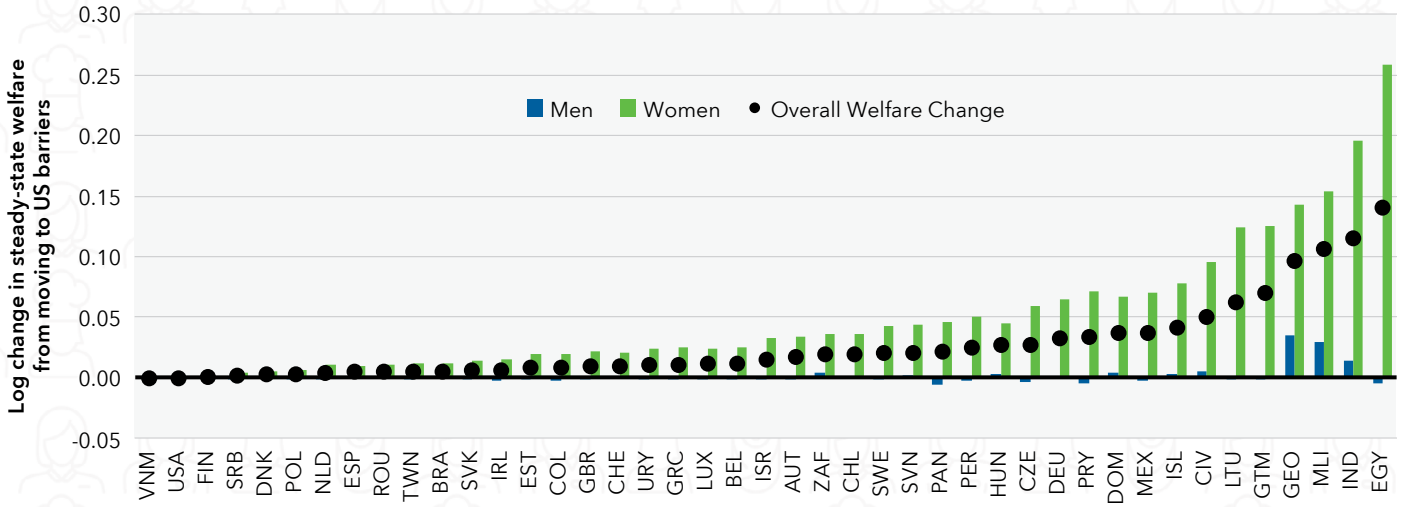

Sources: Li, Zhang, and Zymek, forthcoming. The source of all data for the computation of barriers and welfare counterfactuals is the Luxembourg Income Study database.
Note: Data labels in the figure use International Organization for Standardization (ISO) country codes.

As more women shift into market work, the supply of home services declines, and the welfare effects account for this decline, quantifying the net social gain.

Policy Implications and Conclusions

Our research examines the macroeconomic implications of occupational barriers across various economies that restrict women's labor market opportunities relative to men. We construct a theoretically founded measure that can be used to compare the magnitude of these barriers across countries. We find substantial heterogeneity in occupational barriers for women, with some EMDEs exhibiting barriers comparable to those observed in the US during the 1960s. We show that reducing these barriers could lead to double-digit gains in incomes and welfare for some major emerging markets. The gains arise primarily from improved talent allocation, rather than mere increases in labor force participation.

For economies seeking to accelerate growth in the face of demographic headwinds, these results underscore the importance of policies that effectively reduce these barriers and promote gender equality. The policy solutions will necessarily depend on the context. They could include education and legal reforms, as well as public support for families and childcare. At a time when global growth has slowed and aging populations threaten to shrink workforces worldwide, our findings suggest that removing female labor market barriers represents a powerful lever for boosting productivity and maintaining economic dynamism. For EMDEs in particular, addressing occupational segregation is not just a matter of equity, but also an opportunity to tap a hidden engine of growth that could reinforce their economic development trajectory.

REFERENCES

Hsieh, C. T., E. Hurst, C. I. Jones, and P. Klenow. 2019. "The Allocation of Talent and US Economic Growth." *Econometrica* 87 (5): 1439-74.

Li, N., L. Zhang, and R. Zymek. Forthcoming, "The Allocation of Talent and Income Differences across Countries," Working paper.

United Nations Development Programme (UNDP). 2015. *Human Development Report 2015: Work for Human Development*. Technical Report, New York, NY.

World Bank. 2016. *Women, Business and the Law 2016: Getting to Equal*. Technical Report, Washington, DC.

World Economic Forum (WEF). 2015. *The Global Gender Gap Report 2015*. Technical Insight Report, Geneva.

R&D Misallocation *and the* Productivity Growth Slowdown

Nils Haakon Lehr
nlehr@IMF.org

Illustration: Lanny Nguyen

The Productivity Paradox

Ever since Robert Solow's pioneering work in the 1960s, economists have recognized that long-term economic growth depends primarily on rising total factor productivity (TFP)—how efficiently an economy transforms inputs into outputs. Innovation driven by R&D investment has long been considered the primary engine of TFP growth in advanced economies.

This link appears to have weakened for the US in recent years (Figure 1). While US companies continue to invest heavily in R&D—maintaining a steady 3 percent of GDP over recent decades—the payoff from these investments appears to be shrinking. TFP growth fell from 0.5 percent annually during 1975–95 to just 0.3 percent during 2005–18. This disconnect raises a fundamental question: If we're spending as much as ever on R&D, why is productivity growth slowing down?

Figure 1. R&D Investment and Growth

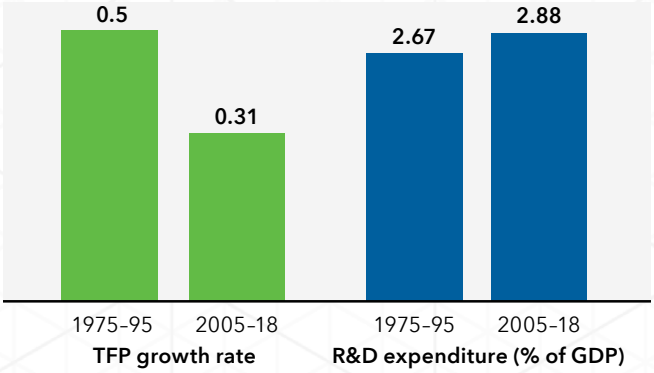

Sources: Penn World Tables 10.01; US Bureau of Economic Analysis (BEA); and Lehr 2025a.

Note: Growth rate of TFP at constant national prices calculated using data from Penn World Tables 10.01. R&D expenditure as a share of GDP calculated using data from the BEA. TFP = total factor productivity.

In two recent studies, we point to an underappreciated explanation: misallocation (Lehr 2025a, 2025b). The problem may not be insufficient investment in innovation, but rather poor distribution of that investment across firms. Think of it this way: If venture capital flowed primarily to mediocre start-ups while breakthrough companies struggled for funding, or if the most talented engineers were stuck in companies where their skills were underutilized, the total amount invested might look healthy, but the outcomes would disappoint.

This is essentially what appears to be happening across the US R&D landscape: Growing **misallocation** of R&D resources across firms appears to draw those resources away from their most productive use, leading to disappointing results even as investment remains high. What are the driving forces behind this misallocation? While there are likely many factors at play, monopsony power over inventors and rising concentration in R&D investments seems to be important forces. Large firms hold back on hiring inventors to keep wages low and thereby push them into less fruitful employment opportunities.

Measuring the Efficiency of R&D Investment

To understand R&D misallocation, consider a simple principle from economics: In well-functioning markets, similar investments should generate similar returns. If one firm earns a 20 percent return on its R&D investment while another earns only 5 percent, market forces should redirect R&D investment, including inventors, toward the more profitable firm until returns equalize. Large disparities suggest that some firms invest too much relative to their innovation capabilities while others invest too little.

Measuring R&D returns is difficult, since it requires quantifying the value created from R&D expenditure by a firm. We propose to measure this value using the estimated valuations of a firm's new patents (Lehr 2025a). Patents are particularly useful here because they are the direct output of innovation activity and because firms are highly motivated to apply for them to protect their intellectual property.[1] Using these valuations, we calculate the R&D return as the value created per R&D dollar for each firm. Think of this as measuring innovation spending's bang for the buck, where patent valuations measure the value innovations add to firms and R&D expenditure the associated costs.

1 As discussed in Lehr (2025a), patents remain an imperfect measure, since not all inventions are patented, for example, because a firm may rely partly on trade secrets to safeguard its intellectual property. We confirm, to the extent possible, that these imperfections are small relative to the documented findings.

The measured R&D returns suggest a striking variance across firms (Figure 2). A firm at the 75th percentile of the R&D return distribution generates nearly twice the innovation value per R&D dollar as a firm at the 25th percentile—and these differences persist for years.

Figure 2. Distribution of R&D Returns across Firms

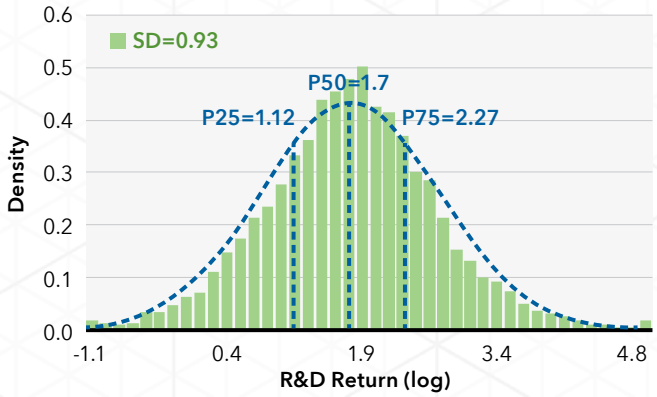

Source: Lehr 2025a.
Note: A histogram showing the highly dispersed distribution of R&D returns, with annotations showing that firms at the 75th percentile achieve nearly twice the returns of median firms.

The Macroeconomic Cost of Getting the Allocation Wrong

Why should policymakers care about these firm-level disparities? Because when R&D resources are misallocated, productivity growth in the entire economy suffers. We develop a framework to quantify this aggregate impact, measuring how much faster productivity could grow if R&D resources flowed to their most productive uses. The key thought experiment is to compare TFP growth under the actual allocation of R&D resources in the economy with a hypothetical allocation in which resources are distributed across firms such that R&D returns are equalized.

Key to this methodology is to distinguish between the private return earned by a firm and the productivity impact, since these can in principle diverge. A long tradition in growth economics emphasizes that the private and public value of an invention can diverge, for example, because some firms are better at extracting economic rent from a given idea thanks to higher markups (Aghion and others 2024; de Ridder 2024). Lehr (2025a) develops a methodology to adjust returns for systematic divergences of public and private impact of innovation using alternative measures of patent quality, including patent citations, to investigate whether firms with high private returns are able to extract a larger share of the value created.[2] Once adjusted, the return can then be aggregated to estimate R&D resource efficiency, the ratio of actual productivity growth relative to its potential—that is, the growth rate achieved at the first-best R&D resource allocation.

The results are sobering. Poor allocation of R&D resources—that is, meager R&D resource efficiency—reduced US productivity growth by an estimated 18 percent during 1975–2014. Put differently, better allocation alone—without any increase in total R&D spending—could have delivered nearly one-fifth faster productivity growth.

More troubling still, the problem has intensified over time (Figure 3). Comparing the late period (2000–14) to the early period (1975–90), worsening R&D allocation accounts for an 11 percent reduction in growth—explaining about one-quarter of the overall productivity slowdown.

Figure 3. R&D Allocative Efficiency over Time

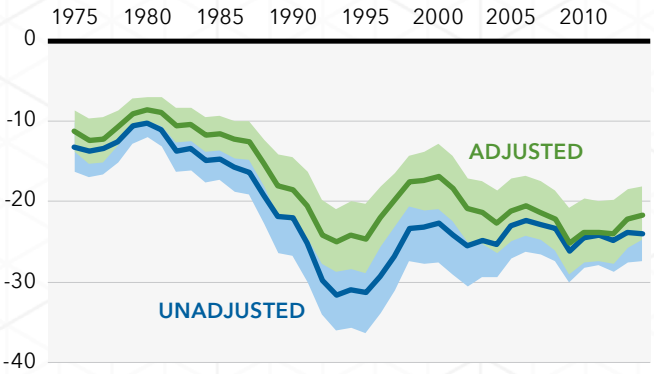

Source: Nils 2025a.
Note: Unadjusted reports R&D allocative efficiency without adjustments for differences in the public and private impact of innovation. Adjusted accounts for such differences.

The Role of Labor Market Power

What prevents more efficient allocation of R&D resources? While there are many potential mechanisms, the data suggest that firms hiring many inventors earn particularly high returns and, thus,

2 The literature commonly uses the number of citations received by a patent as a measure of patent quality or even knowledge externality. This approach assumes that high-quality patents should be cited more, similarly to high-quality academic papers.

appear to be underinvesting. One potential explanation for this pattern is monopsony power in the market for inventors, the focus of Lehr (2025b).

Monopsony power arises when employers can depress wages paid to their employees by hiring fewer of them. To understand monopsony, consider a classic example: a mining town with only one employer. Workers face a stark choice—accept lower wages or relocate entirely. This gives the mining company power to suppress wages below competitive levels, exchanging the possibility of some workers leaving town for the gains from paying the remaining workers lower wages.

Modern inventor labor markets can exhibit similar dynamics. R&D activity is highly concentrated, and many inventors tend to have highly specialized skills—expertise in particular technologies specific to a small set of potential employers. This specialization naturally limits their outside options and, thus, increases the monopsony power of employers.[3] Firms with significant monopsony power over researchers—for example, because they dominate their R&D-specific market—benefit from the ability to suppress wages and, thus, R&D cost. This gives them the high R&D returns documented in the data.

Evidence of Monopsony in Innovation

In Lehr (2025b), we investigate the degree of monopsony for inventors by examining how their wages respond when firms try to expand their R&D workforce. In competitive markets, firms can hire additional workers at the prevailing wage. With monopsony power, expansion requires raising wages to attract workers from other firms. It is exactly this mechanism that gives firms the ability to reduce wages by holding back employment.

Using stock market returns as a source of exogenous variation in firms' demand for inventors, we find clear evidence of monopsony. On average, R&D workers receive only 70 percent of the value they bring to the firm, compared with 100 percent under competitive conditions (Figure 4).[4] But here is the crucial twist: This monopsony power is highly concentrated among firms with large R&D employment. R&D workers at firms with above-median R&D employment receive only 57 percent of the value they contribute, whereas those at firms with below-median R&D employment receive the full value—that is, their compensation is at competitive levels.

Figure 4. Inverse Markdowns by Firm Size

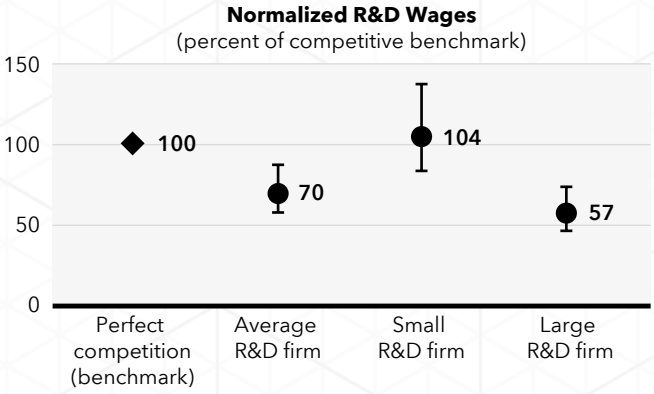

Sources: Lehr 2025b; and IMF staff calculations.
Note: Figure reports inverse markdowns implied by estimated labor supply elasticity. Under perfect competition, wages equal marginal products, implying a markdown of 1.

Quantifying the Innovation Drag

The estimates point to significant size-dependent monopsony, which, in turn, entails significant misallocation: Large, productive firms underemploy inventors to keep their wages low. To assess the full economic impact, in Lehr (2025b), we construct a quantitative model matching the empirical evidence on monopsony power. The results suggest that labor market power in innovation reduces US economic growth by 0.20 percentage point annually—a 13 percent reduction that translates to 11 percent lower welfare.

This occurs through two reinforcing channels: First, there is an **employment effect** as monopsony reduces aggregate inventor employment by about 2 percent. Second, there is an **allocation effect** as size-dependent monopsony shifts inventors toward less productive smaller firms. Remarkably, misallocation accounts for 90 percent of the total impact. Even if monopsony didn't reduce total

3 Monopsony power in this market could also work in the opposite directions. Superstar scientists often possess unique insights and abilities that are crucial for firm success, giving them significant negotiation power. However, most inventors and R&D workers are not superstars and, thus, are more exposed to employer market power.

4 Under competitive conditions, wages equal the marginal product of R&D workers. With imperfect competition, there is a wedge, often referred to as the markdown. The inverse of the markdown is the share of the marginal product that workers receive as wages.

inventor employment, simply distorting the allocation across firms would still cut growth by 0.18 percentage point.

Policy Implications for Reviving Innovation

The findings of these two studies challenge conventional innovation policy focused primarily on increasing R&D spending through tax credits or direct subsidies. While such policies may raise the level of R&D investment, they could worsen allocation if credits flow disproportionately to firms already overinvesting relative to their productivity. Instead, the results suggest that policymakers could achieve significantly more efficient gains through targeted policies that reallocate R&D resources toward a more efficient use.

Instead, the findings of our study suggest that policymakers could consider the following steps:

(1) **Allocation-conscious R&D policy:** Rather than uniform subsidies, efficient use suggests targeting firms with particularly high expected returns on R&D. While efficient targeting might be challenging, given the underlying uncertainty, past returns could be used as an input due to their strong predictive power.

(2) **Strengthening inventor mobility:** Enforcing bans on noncompete agreements, preventing wage-fixing conspiracies, and ensuring portability of benefits could enhance inventors' outside options.[5]

(3) **Antitrust enforcement in innovation markets:** Merger review should consider impacts on inventor labor markets, not just product market competition.

The Path Forward

Our studies deliver a clear message: The US innovation slowdown isn't just about ideas being harder to find—it's about looking for them with an increasingly inefficient system. R&D resources flow to the wrong firms while market power prevents the most productive companies from fully utilizing innovative talent.

The combined welfare costs are substantial. The simulations from our models suggest that R&D misallocation and inventor monopsony together might have reduced welfare by more than 20 percent—this is comparable to other major economic distortions. This suggests that fixing innovation market failures could deliver growth benefits rivaling major tax or regulatory reforms.

For policymakers concerned about productivity and long-term growth, these findings highlight an underappreciated lever: improving how innovation resources are allocated. In an economy where technological progress drives prosperity, ensuring that R&D investments and innovative talent flow to their most productive uses may prove as important as the total amount invested. The challenge now is designing policies that enhance both the level and efficiency of innovation—reconnecting R&D spending with productivity growth.

REFERENCES

Acemoglu, Daron, and Jorn-Steffen Pischke. 1999. "The Structure of Wages and Investment in General Training." *Journal of Political Economy*, 107 (3): 539–72.

Aghion, Philippe, Antonin Bergeaud, Timo Boppart, Peter J. Klenow, and Huiyu Li. 2025b. "Good Rents versus Bad Rents: R&D Misallocation and Growth." NBER Working Paper 34190.

de Ridder, Maarten. 2024. "Market Power and Innovation in the Intangible Economy." *American Economic Review* 1 (114): 199–251.

Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman. 2017. "Technological Innovation, Resource Allocation, and Growth." Quarterly Journal of Economics 132 (2): 665–712.

Lehr, Nils H. 2025a. "Did R&D Misallocation Contribute to Slower Growth?" IMF Working Paper 25/183, International Monetary Fund, Washington, DC.

Lehr, Nils H. 2025b. "Does Monopsony Matter for Innovation?" *Working Paper.*

5 An important caveat to this recommendation is firms' incentives to invest in their workers. As discussed in Acemoglu and Pischke (1999), some restriction on worker mobility can increase firms' incentives to invest in the human capital of their workers.

INDUSTRIALIZATION AND THE BIG PUSH

THEORY AND EVIDENCE FROM SOUTH KOREA

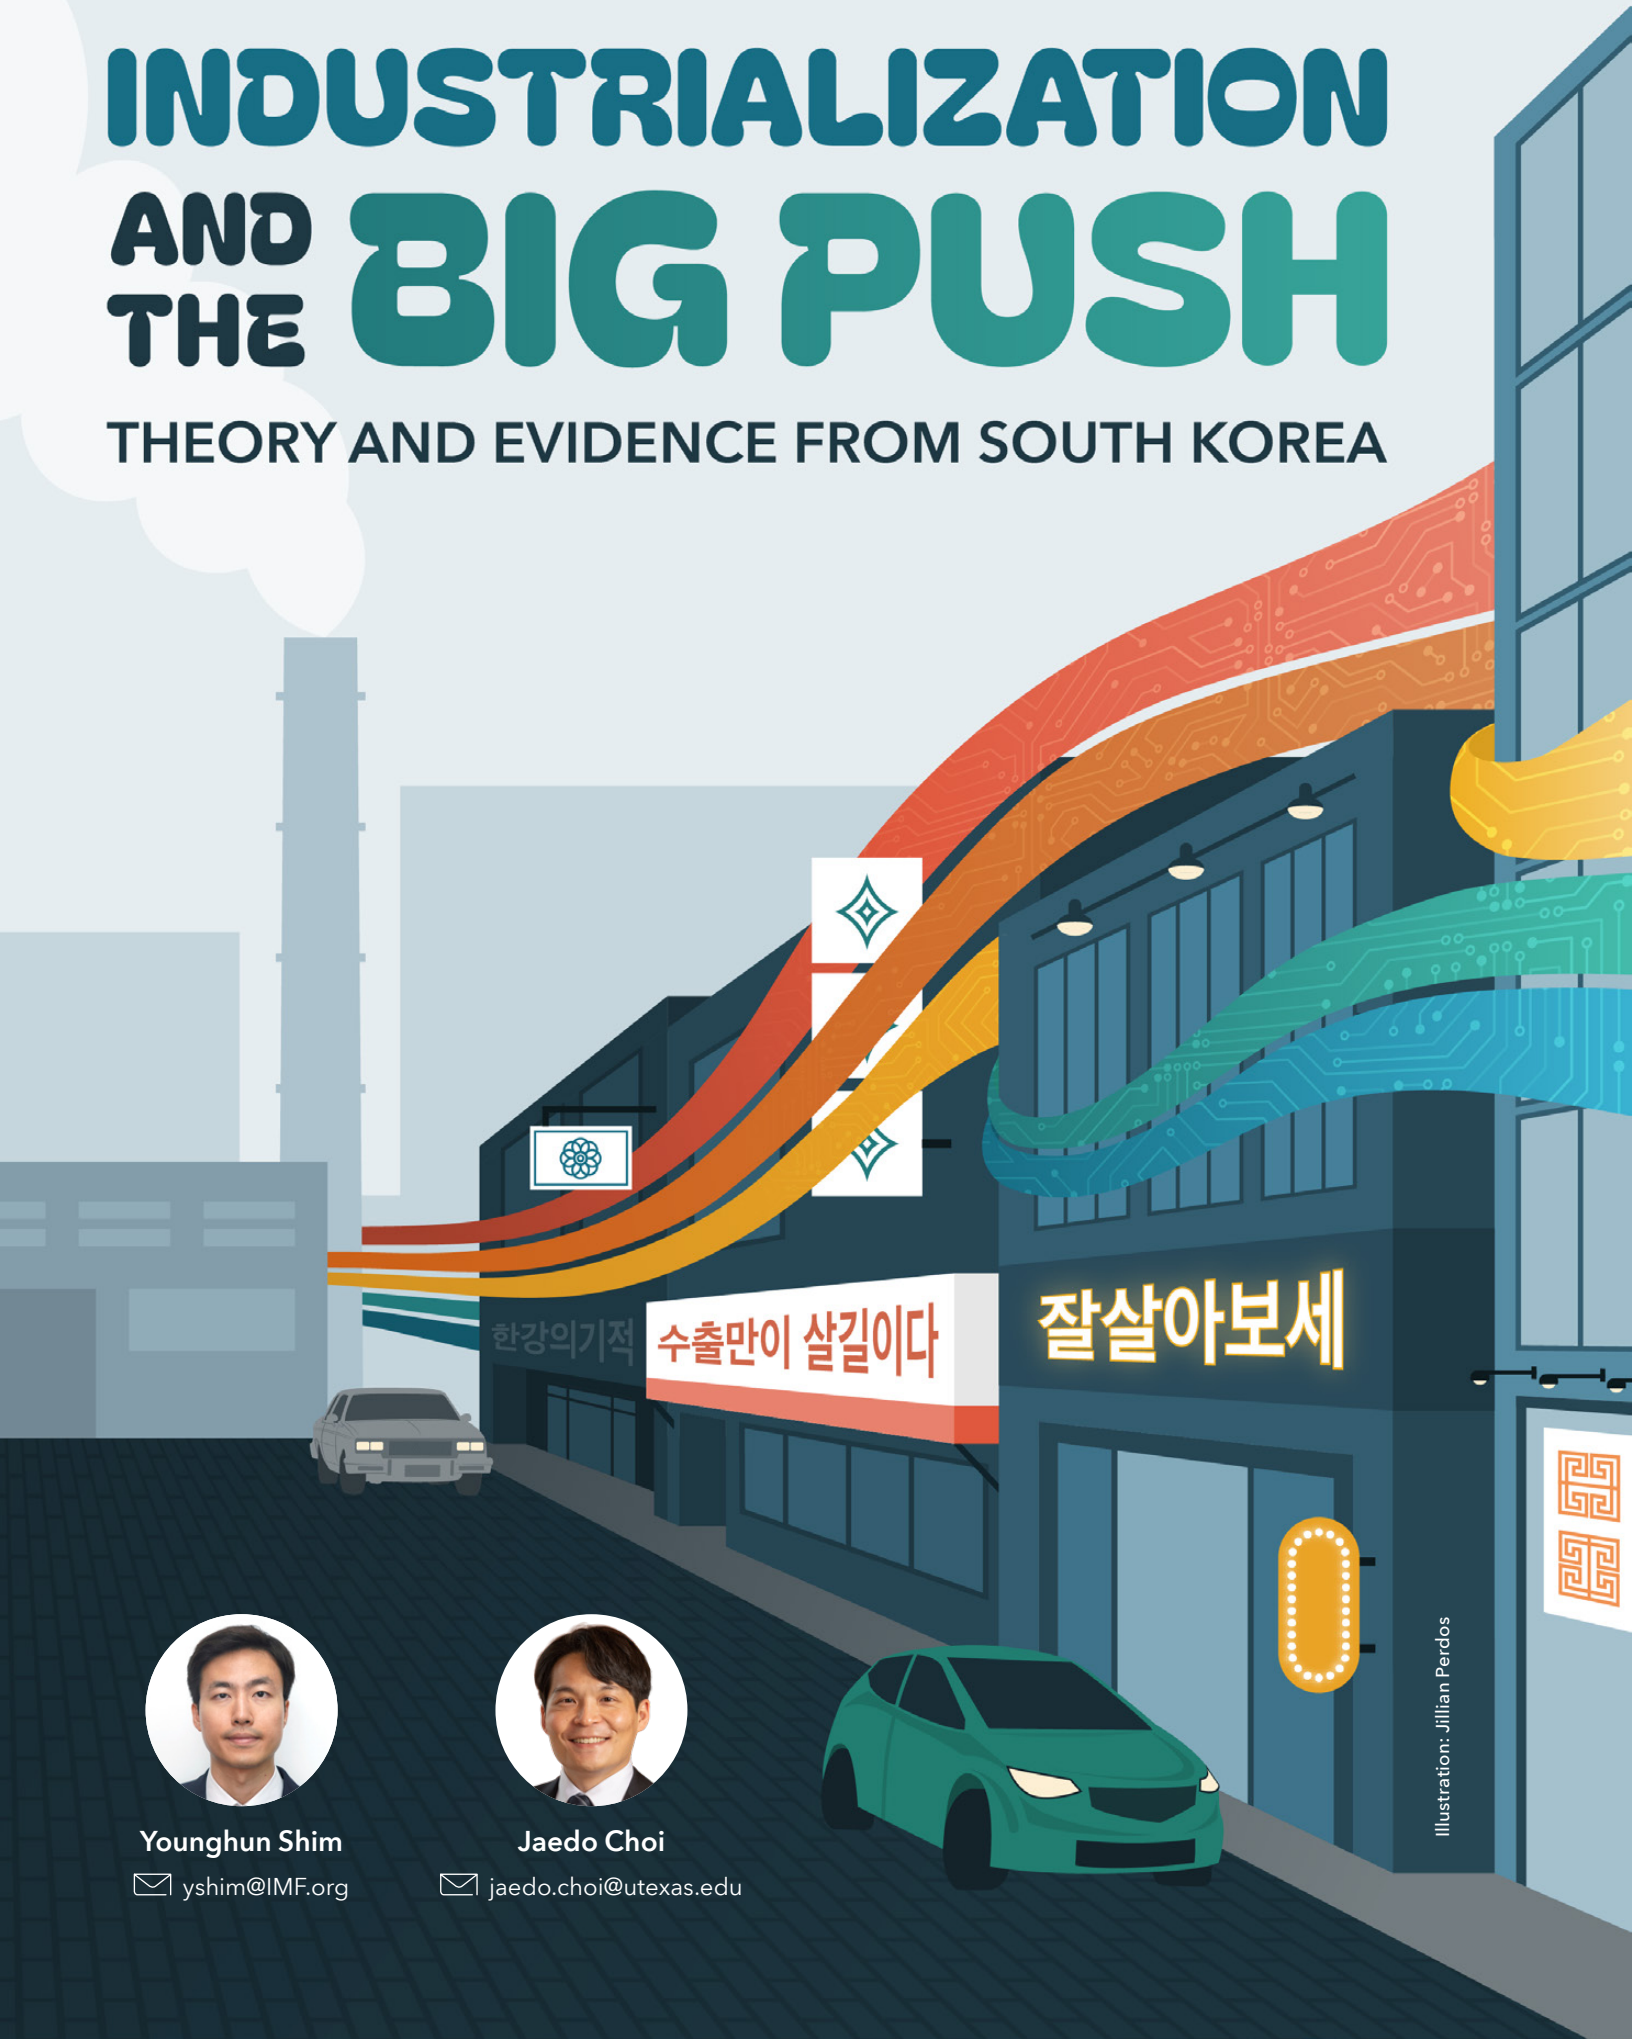

Illustration: Jillian Perdos

Younghun Shim
yshim@IMF.org

Jaedo Choi
jaedo.choi@utexas.edu

Since the seminal work of Rosenstein-Rodan (1943) and Hirschman (1958), economists have argued that *coordination failures* may hinder the adoption of new technologies.

As the benefits of adopting new technologies for any single firm depend on how many *other* firms follow suit, an economy may get stuck in a low-adoption and low-productivity equilibrium, even though coordinating adoption decisions would have made all firms better-off. In theory, one solution is a temporary big push: a coordinated effort to encourage widespread technology adoption by firms, tipping the economy onto a higher-productivity growth path.

Is such a big push possible in practice? We study South Korea's industrial policy in the 1970s that temporarily subsidized technology adoption in heavy manufacturing. Using newly digitized firm-level records of technology contracts with foreign partners, and a quantitative model calibrated to the data, we find strong evidence that the policy succeeded. As predicted by the "big push" hypothesis, the policies helped shift the economy onto a more industrialized, export intensive, and faster long-term growth path. Importantly, the success of the big push was reinforced by complementary measures that improved firms' access to internal and external markets.

South Korea's Heavy-Chemical Industry Drive

Between 1973 and 1979, the Korean government offered temporary support to firms in heavy manufacturing sectors to adopt modern technologies. This period also marked one of the most rapid episodes of industrialization in world history. Our work examines this policy and its long-term economic consequences.

New data: Who adopted what, and when?

As South Korea lagged significantly behind the global technological frontier in the 1970s, Korean firms adopted modern technologies mostly by purchasing technological blueprints from foreign sources.[1] We manually construct a dataset of all technology transfer contracts between Korean and foreign firms during 1970–82, by digitizing the contract documents firms were required to submit to the government. In our context, technology adoption refers to the transfer of *ideas*, such as blueprints or training services, related to mass industrial production.

Our novel dataset reveals a striking pattern for South Korea's industrialization and technology adoption (Figure 1). First, following the policy's launch, the heavy manufacturing share of GDP more than doubled, from 6 percent to 13 percent during the program window, and the number of new technology contracts in these sectors quadrupled. Second, the share of heavy manufacturing in GDP, employment, and exports continued to rise even after the policy ended.

Causal effects of technology adoption

It is challenging to establish a direct causal link between the big push policy and South Korea's unprecedented industrialization. Firms that decided to adopt new technologies might be different from those that did not. Thus, attributing the faster growth of the subsidized sectors to the policy might be misleading. To overcome this challenge, we exploit a winners versus losers research design. Some firms signed technology contracts that were later canceled or delayed by the foreign partner for reasons outside the Korean firm's control (think, for example, of geopolitical shocks or decisions at corporate headquarters abroad). We compare these "loser" firms to otherwise similar "winner" firms whose contracts went forward. Using an event study structure that lines up firms by the timing of (attempted) adoption, we estimate how much technology adoption *in itself* changes outcomes.

To obtain the full picture, we move beyond adopters to study local spillovers and local complementarities. We relate firm growth to changes in the local share of adopters within the same region-sector. To address the possibility that certain places were growing for unrelated reasons, we instrument

1 Using a similar dataset, Choi and Shim (2024b) study the transition from technology adoption to innovation as South Korea caught up to the technological frontier in the 1980s.

Figure 1. Big Push, Adoption of Modern Technology, and Industrialization in South Korea

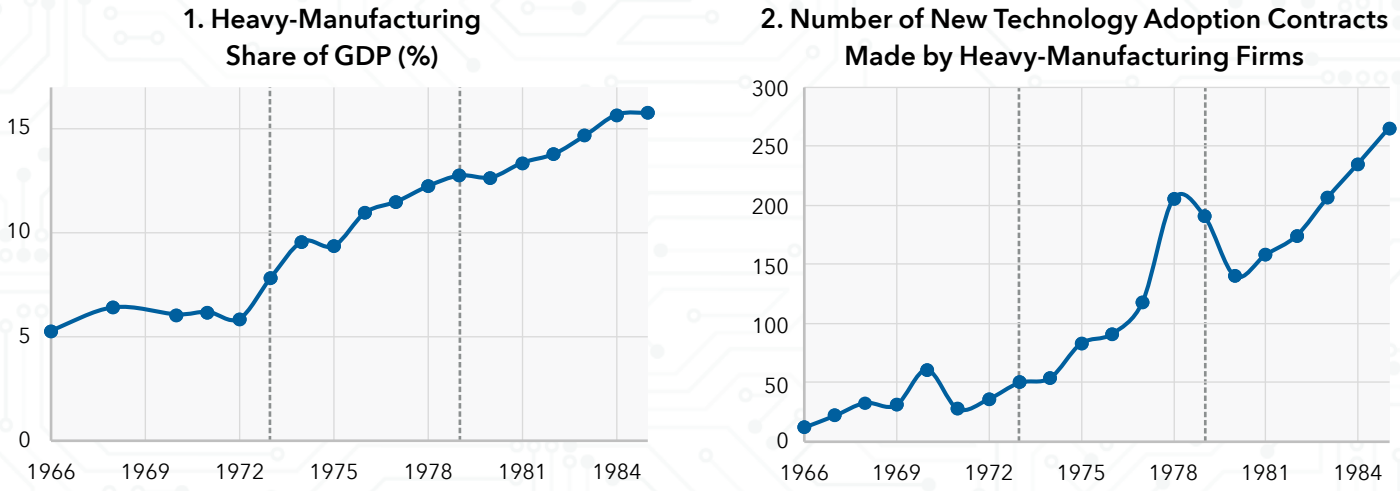

Source: Choi and Shim 2024a.

Note: The vertical lines indicate the start and end of South Korea's big push industrial policy, which subsidized the adoption of modern technologies in heavy manufacturing sectors. Heavy manufacturing sectors include chemicals, electronics, machinery, steel, non-ferrous metal, and transportation equipment. Sectoral value-added data are from the Bank of Korea's input-output tables prior to 1970 and from the Organisation for Economic Co-operation and Development STAN database after 1970. The number of adoption contracts is from the dataset we constructed by digitizing historical archives.

for local adopter shares using the *geography of South Korea's business groups* (chaebols). In particular, we use the fact that technology adoption decisions of firms that are part of the same chaebol are similar even if they are located in different regions, and predict changes in adopter shares inside a region by changes in the business group's adopter share outside the region. We find three main effects.

(1) **Big direct gains for technology adopters.** Firms that successfully adopted modern technologies experienced large improvements: After seven years, their sales increased by 60% and total factor productivity (TFP) were almost double those of observationally similar firms whose contracts fell through (Figure 2).

(2) **Positive local spillovers to nonadopters.** Firms that did not adopt new technologies but were located in regions where technology adoption was higher experienced faster growth in sales and productivity. When the share of local adopters in a region-sector rose by 1 percentage point, nonadopting firms' sales grew by 2.7 percent and their revenue TFP by 1.6 percent. This finding suggests sizable local spillovers, through knowledge diffusion, supplier improvements, and/or demand.

(3) **Local complementarities in adoption.** Adoption itself responded to local adoption: A 1 percentage point increase in the share of local adopters raised a firm's adoption probability by 0.85 percentage point (about 14 percent of the average annual adoption probability in the final year of the program). The complementarity was stronger where the market size was larger.

Together, these results paint a clear picture: Even though private returns to adoption were high for individual firms, adoption was more likely when other firms in a region adopted, and even those that didn't adopt new technologies benefited from other firms' technology adoption. In sum, we find strong evidence of the interdependence that can trap economies in low-adoption equilibrium.

A simple model to connect the dots

To interpret the empirical evidence and quantify the impact of South Korea's big push industrial policy, we build a dynamic model in which firms choose whether to adopt modern technology. Two features of the model are worth highlighting:

- Spillovers occur with a one-period lag: A bigger local base of adopters in period t–1 makes firms more productive and effectively lowers the cost of adoption in period t.

Figure 2. Direct Effects on Adopters: Winners versus Losers Design

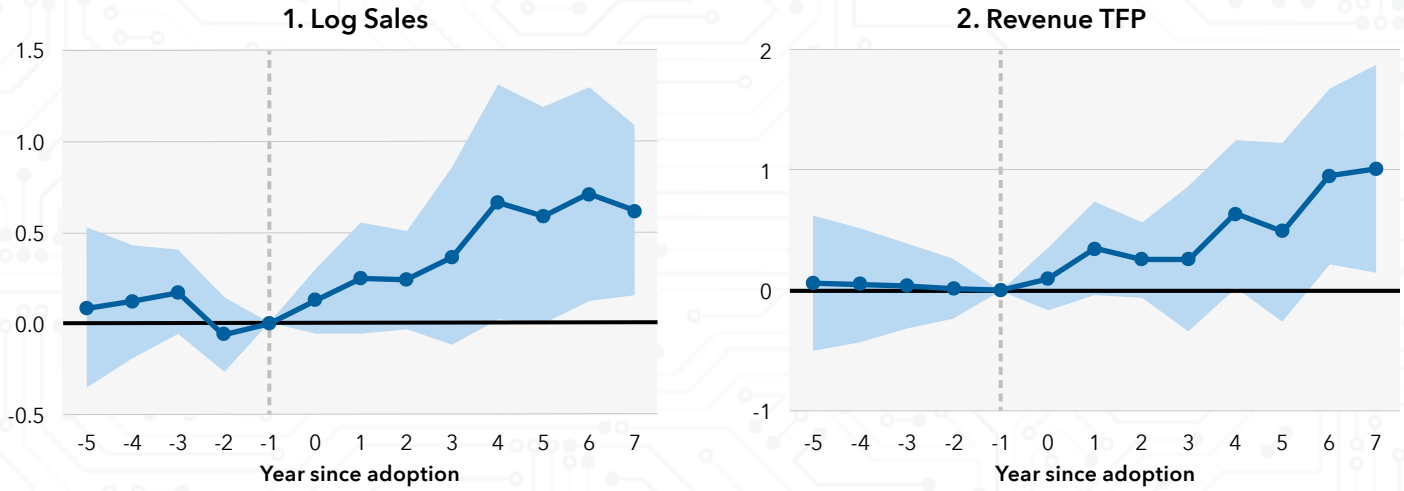

Source: Choi and Shim 2024a.
Note: Panels 1 and 2 present the estimated coefficients of an indicator for adopting new technology on log sales and revenue TFP, respectively, based on the winners versus losers research design. The shaded areas represent the 95 percent confidence intervals based on standard errors clustered at the match and firm levels. All specifications include match-year and match-firm fixed effects. For further details, see Choi and Shim 2024a. TFP = total factor productivity.

- Adoption costs are fixed: The decision to adopt is similar to a one-off investment whose returns are larger the stronger the spillovers and the larger the scale of a firm.

These model ingredients generate dynamic complementarities. More adoption today makes adoption tomorrow more attractive. We show analytically that if spillovers and private returns lie in a *middle* range—neither too weak nor too strong—the model admits multiple steady states: a low adoption "preindustrialized" state and a high adoption "industrialized" state. In that case, a temporary technology adoption subsidy can have a permanent effect by pushing the economy from the low to the high steady state.

Did the 1970s policy really cause a long-term shift?

In addition to the features discussed above, our calibration model also incorporates costly internal and external trade and input-output linkages across industries. We calibrate the model to match our causal estimates of (1) direct gains and local spillovers from technology adoption, which tightly pin down the two key parameters governing private returns and externalities, and (2) cross-regional patterns linking adoption to market access. We then ask: *What would have been Korea's growth trajectory if there had been no technology adoption subsidies?*

Counterfactual result. Without the 1970s big push subsidies, South Korea would have converged to a less industrialized steady state (Figure 3). Our quantitative simulations suggest that in the absence of the policy, the heavy-manufacturing share of GDP and export intensity would have been 27 percent and 39 percent lower, respectively, than in the baseline scenario with the policy in place. The big push policy raised aggregate welfare by 14.6 percent, but gains were uneven across regions, with welfare changes ranging from -1.2 percent to 83.8 percent. These disparities reflect two opposing forces triggered by localized productivity improvements. On one hand, internal trade linkages transmit lower prices, benefiting consumers and input-using firms elsewhere. On the other hand, stronger domestic competition compresses profits in some regions.

Why market access complements a big push

The strength of complementarity in our setting rises with the scale of a firm, which in turn depends on access to international markets and input costs. We therefore examine three "What if?" scenarios in which scale was temporarily reduced during the policy years:

- Foreign demand held flat rather than growing
- Higher import tariffs (raising the cost of imported intermediates)

Figure 3. Aggregate Effects of the Big Push: Baseline versus Counterfactual

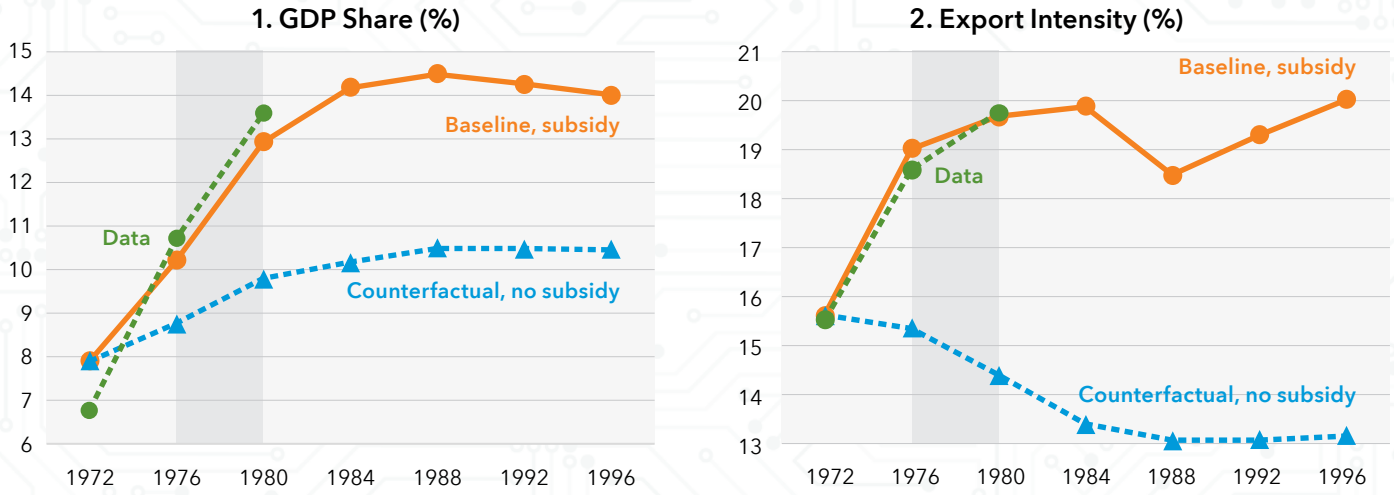

Source: Choi and Shm 2024a.

Note: Panels 1 and 2 show the share of heavy manufacturing in GDP and exports in the baseline and counterfactual economies. The shaded areas indicate the period of South Korea's big push industrial policy. The green lines represent the data from the input-output tables, while the orange and blue lines are the outcomes of the baseline and counterfactual economies, respectively.

- No major highway improvement (higher internal trade costs)

Each scenario weakens the impact of the subsidies; holding down foreign demand is especially powerful. In the combined scenario—flat foreign demand, higher import tariffs, and no highway improvement—the big push does not occur even with the same subsidies. The takeaway is simple: *Market access and complementary policies magnify (or mute) the power of temporary adoption subsidies.*

Conclusion

We empirically and quantitatively explore the possibility of industrialization through a big push for technology adoption, focusing on South Korea's experience in the 1970s. Our findings suggest that a big push policy can be successful by addressing coordination failures and facilitating the diffusion of advanced technologies. In addition, we uncover the importance of complementary policies that ensure sufficient market access for sustainable industrial growth.

REFERENCES

Choi, Jaedo, and Younghun Shim. 2024a. "Industrialization and the Big Push: Theory and Evidence from South Korea." IMF Working Paper 24/259, International Monetary Fund, Washington, DC. https://doi.org/10.5089/9798400292927.001.

Choi, Jaedo, and Younghun Shim. 2024b. "From Adoption to Innovation: State-Dependent Technology Policy in Developing Countries." IMF Working Paper 24/154, International Monetary Fund, Washington, DC. https://doi.org/10.5089/9798400281341.001.

Hirschman, Albert O. 1958. *The Strategy of Economic Development.* New Haven, CT: Yale University Press.

Rosenstein-Rodan, P. N. 1943. "Problems of Industrialisation of Eastern and Southeastern Europe." *The Economic Journal* 53 (210/211): 202–11.

INFORMATION, SOCIAL MEDIA, and INTERNATIONAL TRADE

Theory and Evidence Using 20 Million Online Postings

George Cui
gcui@IMF.org

Kailin Gao
kailingao@uibe.edu.cn

Illustration: Lindsey Kim

A long-standing puzzle in economics is how businesses get the information they need to succeed. This is especially true in international trade, where a company must acquire vast amounts of information from foreign regulations and shipping logistics to the tastes of customers thousands of miles away. The challenges of finding reliable information—what economists call "information frictions"—can prevent a company from ever selling its products abroad.

For decades, researchers were unable to observe how firms obtain information: Interactions between firms were invisible, leaving what Paul Krugman called "no paper trail." (Krugman 1991) But the rise of the internet has changed that. Digital platforms now allow even the smallest firms to find partners and solve problems together across vast distances.

Our research shines a light into this previously hidden world (see Cui and Guo 2025). We analyzed over 20 million online posts on a Chinese social media platform where firms and individuals share practical advice about international trade. We linked these digital conversations to official customs data, which allowed us, for the first time, to measure how this peer-to-peer learning impacts a company's ability to export.

Our findings are striking. We found that joining the platform gives a firm's exports a significant and lasting boost. In other words, the information acquired by firms through social media interactions on the platform impacts their actions. Moreover, we discovered that information flows between firms are shaped by their strategic relationship in supply chains. Upstream firms tend to share information with their downstream partners, hoping that doing so will increase their own sales. However, access to the benefits of the social media platform is not equal: Firms in China's more developed regions are far more likely to participate, potentially creating a new kind of "digital divide."

To understand the macro-relevance of these findings, we construct a quantitative general equilibrium trade model. The model integrates endogenous information acquisition and sharing mechanisms in a monopolistically competitive market structure with heterogeneous firms. Both our real-world data and our model show that social media can be a powerful tool for reducing the information barriers that hinder trade. The policy implications are clear: To help businesses succeed globally, governments should not only focus on cutting trade costs such as tariffs, but on supporting the knowledge networks that empower firms to help themselves.

Online Postings: A New Window on Firm-to-Firm Learning

Our study analyzes data from one of China's largest social media platforms for trade professionals. Founded in 2002, the platform serves as a digital town square for the managers, salespeople, and owners of companies engaged in international commerce. Registration is free, and users can choose to provide their company's name, industry, and location.

The platform is, in essence, a massive, user-driven export knowledge base. To make this concrete, consider an example from our data during the 2018 US-China trade war. One user posted a simple urgent request: "Let's talk about how to avoid taxes (US trade war)!!" The answers were immediate and practical. A fellow user outlined several options, including rerouting goods through a third country for minor processing or moving the factory to a country not subject to the tariffs. This single exchange is a snapshot of how practical, valuable business knowledge is shared in real time. The platform has dozens of subforums on topics such as industry communication, support services, and business communication.

The Digital Dividend: Quantifying the Export Gains

Does all this online talk actually help companies sell more abroad? To answer this, we had to overcome a key challenge: It is possible that companies that are already growing faster are more likely to join

Figure 1. The Impact of Platform Adoption on Firm Exports

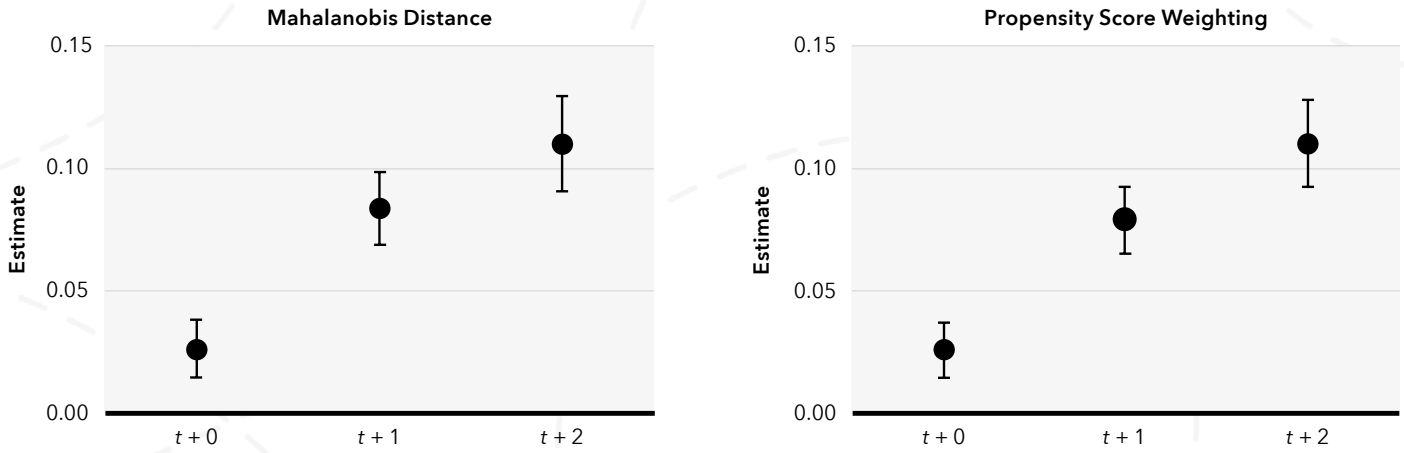

Note: This figure shows the estimated impact of Platform Adoption on firm-level export volume using Mahalanobis matching (left panel) and propensity score (PS) weighting (right panel). The black dots represent point estimates, and the vertical lines indicate 95% confidence intervals. The x-axis denotes years relative to the intervention (t=0 is the year of access). Estimates suggest a positive and increasing effect over time, with statistical significance in most periods.

an information platform. If so, simply comparing the sales of platform members with nonmembers would be misleading.

To gauge the platform's effect on firm exports, we used a sophisticated statistical technique that carefully creates a valid control group. The method matches each company that joins the platform with nearly identical companies that haven't joined yet, based on their industry, size, export history, and other characteristics. We ran multiple checks to confirm that these groups were on the same trajectory *before* the treatment groups joined the platform. This suggests that any later difference in performance is likely a result of the knowledge gained from the platform itself.

The gains from information sharing on the platform are remarkable. In the first year after a company's employee joins the platform, its exports increase by an average of 3 percent. The benefit grows significantly over time, with exports climbing by 8 percent in the second year and 11 percent by the third year. This rising effect suggests that companies don't just get a one-time boost; they continue learning how to use the network more effectively over time, leading to lasting gains.

The benefit of joining the social media platform is also amplified by the size of the network. We found that companies in regions with a larger existing community of active users gain even more when they join. This points to what economists call "external economies of scale"—the idea that a company benefits when other firms in its area are also active and sharing knowledge. A 10 percent increase in the local pool of information is linked to an extra 0.3 percent in export growth for a new member. The platform isn't just a static library of information; it's a dynamic, self-reinforcing ecosystem where each new question, answer, or shared experience enriches the collective knowledge base. As more firms join and contribute, the value of the network grows—reducing barriers to entry, accelerating learning, and creating a virtuous cycle of growth that benefits even those that have yet to participate.

The New Map of Knowledge: How Information Really Flows

Knowing that the platform works, we next investigated *how* information moves between platform users. Our findings revealed two patterns that challenge common assumptions about the digital age.

First, geography still matters. In a world of instant communication, one might think that distance is no longer a barrier. Our data dispute this view: After controlling for other factors, we found that a 10 percent increase in the physical distance between two cities reduces the volume of online messages between firms in those locations by about 0.7 percent. This leads to a new puzzle: If it costs nothing extra to reply to a user across the country versus to one across the street, why does distance still matter for information sharing?

Figure 2. Knowledge Flows Are Strategic, Not Random

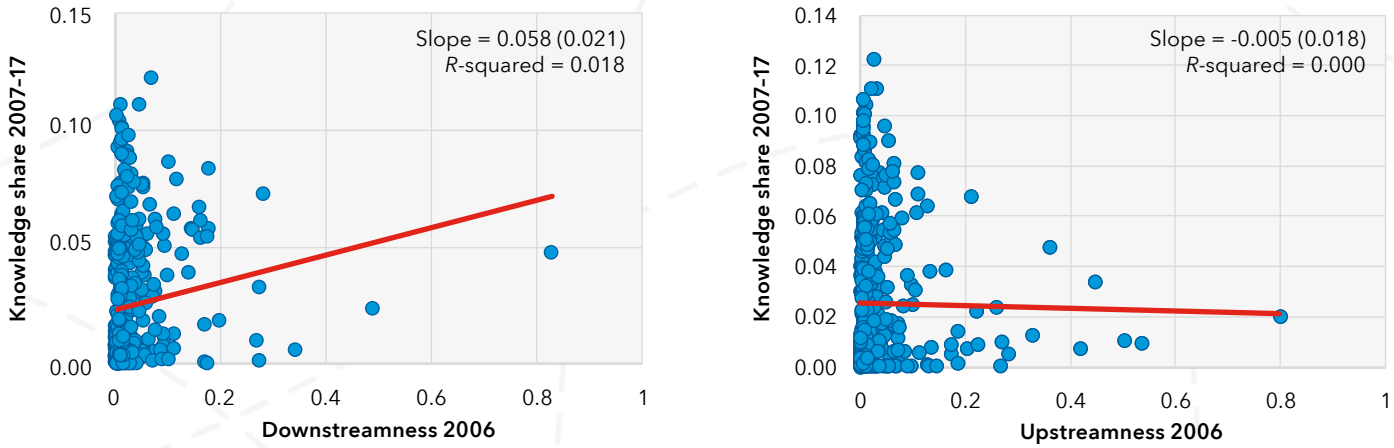

Note: This figure displays the scatterplots of the knowledge shares between different sectors on the y-axis against the relative downstreamness (left) and the relative upstreamness (right). The "knowledge share" is the share of teaching sector's knowledge outflow (measured by posting number) to the learning sector divided by the teaching sector's total knowledge outflow from 2007 to 2017 on the platform. The relative downstreamness of the learning sector to the teaching sector is measured by the share of teaching sector's total input production purchased by the learning sector in 2006 WIOD. The relative upstreamness of the learning sector to the teaching sector is measured by the share of teaching sector's total input purchase supplied by the learning sector in 2006 WIOD. All panels exclude within sector observations.

The solution to this puzzle lies in our second finding: Information is shared for strategic reasons, and it flows through supply chains. We discovered that knowledge is far more likely to be shared by an upstream supplier with its downstream customer compared with other platform users. In other words, companies use the platform to educate their buyers, therefore improving the coordination and efficiency of the supply chain. A sector is much more likely to share information with another sector if the second sector is a major customer. When a buyer's purchases from a supplier increase by 10 percent, the supplier, in turn, shares about 0.6 percent more knowledge with that buyer.

This "learning from sellers" dynamic makes perfect business sense. A company that produces electronic components has a direct financial interest in successful exports by the downstream firms that assemble those components into manufacturing goods. A successful buyer means more future orders for the supplier. This explains why distance still matters: A company is more likely to have these important supplier-customer relationships with firms that are physically closer. This finding upends the traditional view of knowledge diffusion that says good ideas flow mostly at random. Instead, we show that firms treat knowledge sharing as a deliberate, strategic investment in their own business network.

Why does strategic knowledge diffusion matter for policymakers? Understanding that firms share knowledge for strategic reasons—especially along supply chains—offers policymakers a roadmap for smarter, more targeted interventions to spur growth. Rather than relying solely on broad export-promotion campaigns or generic digital infrastructure investments, governments can leverage firms' existing natural incentives.

First, our findings suggest that not all firms need to be subsidized equally. Suppliers that stand to gain the most from their customers' export success are already motivated to share what they know—but they may lack the tools or proper channels to do so. Policymakers can support this dynamic by investing in platforms that make peer-to-peer knowledge sharing easier, more trusted, and more discoverable. Think of it as building digital highways that connect input producers with final exporters.

Second, the government could help design platforms that reward knowledge sharing, not just consumption. If knowledge sharing is strategic, platforms should be built to reflect that. Public or semipublic export support platforms could include features that recognize, rate, and reward valuable contributions—such as verified answers, case studies, and tutorials. Just as e-commerce platforms reward top sellers, trade forums can be designed to highlight top mentors, nudging firms to invest more in sharing.

Figure 3. The Digital Adoption Gap

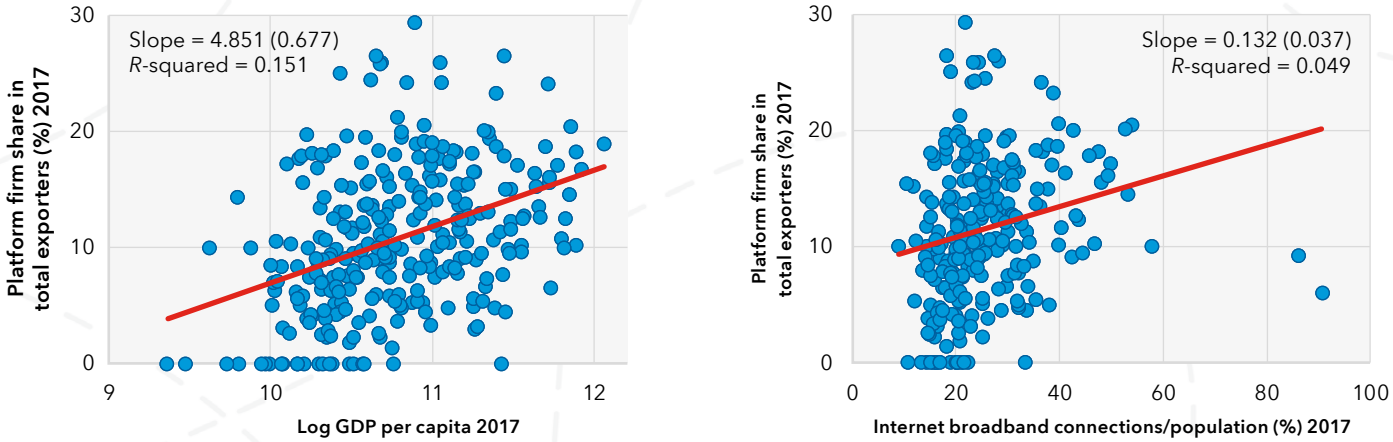

Note: This figure displays the scatterplots of platform firm share in total exporters against the city-level logarithm GDP per capita and internet access.

Third, understanding that distance still shapes knowledge flows—even in digital space—means that closing geographic gaps remains essential. Strategic partnerships, "twinning" programs between inland and coastal cities, and hybrid online-offline forums can help bridge these divides. Without such efforts, the knowledge-rich get richer, and lagging regions fall further behind.

In summary, when knowledge is treated not as a public good that flows randomly, but as a strategically traded asset, policy can be designed to support and amplify those incentives. The result: faster diffusion of know-how, stronger export ecosystems, and more inclusive growth.

The Digital Divide and the Role of Policy

Although the platform is a powerful tool, not everyone has equal access. Our research uncovered a digital divide: Cities with higher GDP per capita and better internet infrastructure have higher rates of adoption of the digital platform.

This could create a troubling dynamic. The regions that are less developed and arguably are in need of more information are less likely to get it, which could worsen regional inequality over time. This could occur because firms in less developed areas face higher barriers to getting online and, once there, may find a less robust network to help them.

This finding suggests that providing support for the development of and access to digital platforms in lagging regions may help attenuate the digital divide and regional inequality.

Modeling a World with Digital Knowledge

To understand the macroeconomic importance of these findings and to evaluate alternative policies, we built a general equilibrium trade model, in which firms choose both to learn and to share information.

In our model, the economy has multiple regions and industries that are connected through supply chains. Firms have different levels of efficiency and can choose to pay a "learning fee" to access a shared pool of knowledge. Tapping into this knowledge may give firms a new "idea" that makes it cheaper and easier to export. The decision to *share* information is also a key part of the model. Firms behave strategically and share information in a way that maximizes the profits of their own local business networks.

We calibrate the parameters of the model to match the patterns in our data. For instance, the export boost from joining the platform determined the "value of an idea" in our model. The tight link between data and theory strengthens the reliability of the model's predictions.

Conclusion: Establish Information Networks

The rise of social media has changed the way people connect and share ideas. Our research shows that for businesses engaged in international trade, these platforms are becoming an essential digital economic infrastructure. By allowing

companies to learn directly from their peers, these networks reduce information barriers, leading to real, sustained increases in exports.

Our work offers three key lessons. First, the spread of business knowledge is not random; it is driven by clear economic incentives within supply chains. Second, even in our hyperconnected world, geography and local business relationships still matter. Third, unequal access to these digital tools can widen the gap between prosperous and lagging regions.

These insights point to new directions for trade policy. To help companies compete globally, governments can do more than just cut tariffs. A complementary strategy may be to invest in the digital infrastructure that allows firms to build knowledge networks for themselves. Supporting access to these platforms or to a relevant knowledge base, especially for businesses in underserved areas, could be a high-impact, low-cost way to foster more inclusive economic growth. In the 21st century, a smart trade policy must also be a smart information policy.

REFERENCES

Cui, George, and Kailin Gao. 2025. "Information, Social Media and International Trade: Theory and Evidence Using Twenty Million Online Postings." Working paper.

Krugman, Paul. *Geography and Trade*. 1992. Cambridge, MA: MIT press, 1991.

THE EVOLVING LANDSCAPE OF GLOBAL TRADE AND FINANCIAL INTEGRATION

The IMF Research Department held the 26th Jacques Polak Annual Research Conference, November 6–7, 2025.

This year's conference focused on "The Evolving Landscape of Global Trade and Financial Integration." The conference provided a forum to discuss innovation research and facilitated exchange of ideas among researchers and policymakers.

In the policy panel titled "A World in Transition: Are We Ready to Adapt?," Hélène Rey, Kristin Forbes, Philip Lane, Sethaput Suthiwartnarueput, and Pierre-Olivier Gourinchas discussed the main challenges confronting a world undergoing significant shifts. They shared their perspectives on whether economies are prepared to meet these challenges and how policymakers can best adapt.

The Mundell-Fleming Lecture was delivered by Ricardo Reis, who analyzed the role of financial repression in the 21st century. He presented a framework to think through financial repression and provided estimates of the revenues these measures can generate. While repression revenues are likely to average zero over the long run, Reis showed that they can be sizable over a decade or two, helping explain why policymakers may resort to them despite their potential distortions.

Jiaqian Chen
JChen@IMF.org

The joint IMF-FCDO Conference marked 13 years of collaboration between the IMF and the United Kingdom's Foreign, Commonwealth & Development Office (FCDO) on the Macroeconomic Research in Low-Income Countries (MRLIC) program.

The event brought together researchers, policymakers, and development partners to take stock of the achievements of the long-standing strategic partnership and discuss the challenges faced by low-income countries (LICs) in a rapidly changing global landscape.

The discussion focused on issues such as geoeconomic fragmentation, demographic transitions, rapid technological change, and extreme weather events as they relate to LICs. The conference featured panel discussions and pathfinding sessions on building resilience, fostering inclusive growth, and leveraging data and innovation to improve policymaking.

The conference outlined areas for future research and highlighted the need for partnerships to support macroeconomic research on LICs. IMF Managing Director Kristalina Georgieva and FCDO chief economist Dennis Novy gave the opening remarks; IMF Deputy Managing Director Nigel Clarke and IMF UK Executive Director Veda Poon highlighted the main takeaways from the conference in closing.

Petia Topalova
ptopalova@IMF.org

IMF RESEARCH
perspectives

Learn more about IMF research at
IMF.org/en/research

Visit *IMF Economic Review*,
the official research journal of the IMF, at
https://link.springer.com/journal/41308

Find IMF publications by visiting
IMF.org/pubs

View IMF research and data for free at
eLibrary.IMF.org