NBER | NATIONAL BUREAU of ECONOMIC RESEARCH

The Reporter

No. 1, 2024

In This Issue

Program Report: International Finance and Macroeconomics

Mark A. Aguiar and Linda L. Tesar

Affiliates of the International Finance and Macroeconomics (IFM) Program study financial interactions among nations, including cross-border capital flows, exchange rates, responses to global financial crises, and the transmission of economic shocks. Rather than attempting to summarize the more than 1,000 working papers these researchers have distributed since the last program report in 2015, we focus here on three issues that have attracted substantial research attention from this group: impacts of the COVID-19 pandemic, global supply chain shocks, and the privileged position of the US dollar in global asset markets.

Impact of the COVID-19 Pandemic

The COVID-19 pandemic was a large, global shock with long-lasting economic impact. About a quarter of the IFM working papers submitted since 2016 relate directly to this shock and its aftermath. Figure 1 illustrates the magnitude of the COVID-19 shock around the globe. The impact on the United States was dramatic: a 10 percent decline in GDP between the second quarter of 2019 and the second quarter of 2020. Peru, India, Spain, and the United Kingdom suffered declines in economic activity in excess of 20 percent.[1]

Economists deployed a wide range of tools to try to make sense of both the magnitude of the decline in economic activity as well as its transmission through the population and across national borders. One approach, illustrated in research by Zachary Bethune and Anton Korinek; Anand Chopra, Michael

Figure 1

GDP Declines during COVID-19

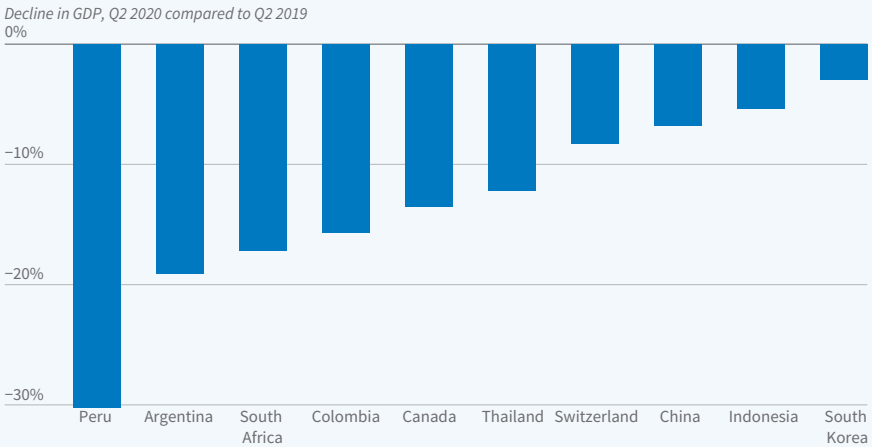

Source: Enrique G. Mendoza, Eugenio I. Rojas, Linda L. Tesar, and Jing Zhang. "A Macroeconomic Model of Healthcare Saturation, Inequality, and the Output-Pandemia Trade-off," NBER Working Paper 28247 and in IMF Economic Review 71, November 2022, pp. 243–299.

Research Associates Mark A. Aguiar and Linda Tesar are the codirectors of the NBER's International Finance and Macroeconomics Program. Aguiar is the Walker Professor of Economics and International Finance at Princeton University. Tesar is the Alan V. Deardorff Collegiate Professor of Economics at the University of Michigan.

Devereux, and Amartya Lahiri; and Alexander Chudik, M. Hashem Pesaran, and Alessandro Rebucci, incorporates the Susceptible, Infected, and Recovered framework of viral transmission into a macroeconomic model that highlights the trade-off between lockdowns to limit contagion and the consequent decrease in income.[2] Workplace interactions cause the disease to spread, illness reduces productivity and causes workers to retreat from the workplace, and, ultimately, high rates of infection prompt governments to impose policies on social distancing.

Taking a slightly different perspective, Enrique G. Mendoza, Eugenio I. Rojas, Tesar, and Jing Zhang use international evidence to show that saturation of the healthcare systems — rather than high infection rates — trigger costly economic lockdowns.[3] They document that shutdowns happen even when infection rates are relatively low due to bottlenecks in the healthcare sector. In the acute phase of the crisis, lower-income countries were less likely to impose strict lockdowns and therefore experienced larger mortality rates. Cristina Arellano, Yan Bai, and Gabriel P. Mihalache find that countries with less borrowing capacity face tighter fiscal constraints, making lockdowns more costly and constraining policymakers' ability to impose lockdowns and limit the spread of the disease.[4] Laura Alfaro, Oscar Becerra, and Marcela Eslava examine the transmission dynamics of COVID-19 in emerging markets, where a greater share of economic activity is conducted by small, informal firms that cannot pivot to working remotely.[5]

At the outbreak of the pandemic, uncertainty over the severity and duration of the crisis triggered a sudden increase in excess demand for liquid assets and a sudden drop in forecasts of future economic activity. Banks and nonbank financial institutions increased their holdings of precautionary cash balances; at the same time, firms drew down their bank lines of credit, generating a spike in risk premia and a fall in equity returns.[6] Linda S. Goldberg and Fabiola Ravazzolo, and Michael Bordo, document how, in response to this dash for cash, the Fed

injected liquidity by providing swap lines and central banks worked in concert to prevent a global financial crisis.[7] The flight to safety in times of uncertainty underscores the critical role of the dollar during times of volatility, a topic that we will return to below.

Policy during COVID

Provision of liquidity was just the first of many policy measures undertaken to support the fragile global economy. To offset the negative impact of the pandemic, governments provided loans to firms, made transfer payments to households, expanded unemployment coverage, intervened in foreign exchange markets, relaxed macroprudential regulations on financial institutions, and lowered interest rates. Though the set of policy instruments deployed was similar across countries, the precise policy mix varied widely, as pointed out by Joshua Aizenman, Yothin Jinjarak, Hien Nguyen, and Ilan Noy.[8] Studying a large cross-section of countries, Katharina Bergant and Kristin Forbes find that the choice of policy instrument is strongly influenced by the "policy space" available for that particular tool.[9] For example, countries with low interest rates entering the pandemic had less capacity to lower rates further; countries with lax financial regulations had less capacity to loosen regulations. Interestingly, a lack of fiscal space did not constrain advanced economies from using fiscal policy despite their entering the pandemic with high levels of sovereign debt.

Figure 2, from a study by Pierre-Olivier Gourinchas, Şebnem Kalemli-Özcan, Veronika Penciakova, and Nick Sander, shows the increase in fiscal expenditures as a share of GDP in 2020.[10] The increases in spending were unprecedented for peacetime governments — an average 16.6 percentage point increase in advanced economies. In addition, there were large increases in "below the line" loans and guarantees to households and firms in much of Europe. Fiscal spending was more modest in emerging markets, where it was limited by the degree to which those governments could engage in debt-financed spending.

Figure 2a

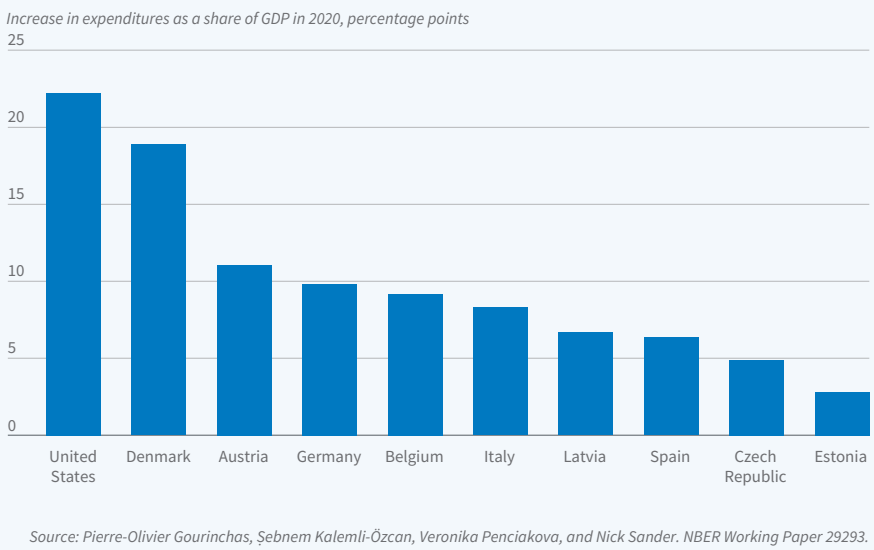

Source: Pierre-Olivier Gourinchas, Şebnem Kalemli-Özcan, Veronika Penciakova, and Nick Sander. NBER Working Paper 29293.

Figure 2b

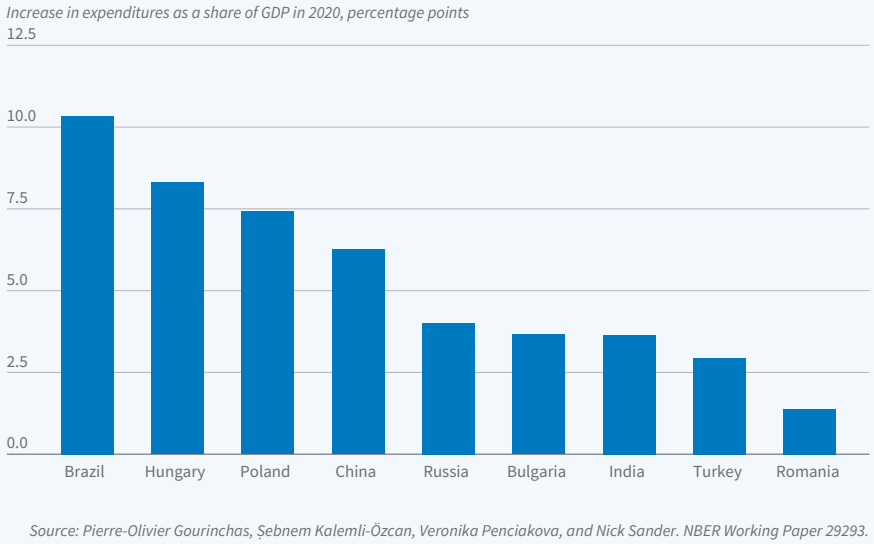

Source: Pierre-Olivier Gourinchas, Şebnem Kalemli-Özcan, Veronika Penciakova, and Nick Sander. NBER Working Paper 29293.

Did this unprecedented level of spending help support struggling firms? Gourinchas et al. drill down to examine the impact of fiscal policy on small and medium-sized enterprises across 27 countries and 50 sectors. Among their many findings, they report that fiscal policy succeeded in preventing widespread failures of small and medium-sized firms, but was less successful in targeting the firms that needed the most help. Fiscal expansion in advanced economies had small or even negative spillovers to emerging markets, and therefore did not make up for the smaller fiscal packages in emerging markets.

Global Supply Chains and Inflation

The pandemic underscored the world's vulnerability to supply chain disruptions. Tensions between the United States and China prior to the pandemic had already induced many multinationals to reorient their supply chains toward trading partners deemed friendlier to US interests or to

on-shore their activities. Lockdowns and plant closures in China during the pandemic put pressure on an already strained global supply network. Bottlenecks in transportation and at ports of entry delayed delivery of final goods as well as intermediate inputs used in manufacturing. These disruptions led some policymakers to question the benefits of trade and to push for further renationalization of supply chains. Barthélémy Bonadio, Zhen Huo, Andrei A. Levchenko, and Nitya Pandalai-Nayar quantify the role of supply chains during COVID using a general equilibrium framework that captures the input-output linkages between industries and across borders.[11] They find that about a quarter of the decline in economic activity during COVID can be attributed to global supply chain disruptions. Further, counterfactual exercises suggest that a renationalization of supply chains, reducing reliance on international product flows, would only exacerbate the output decline. Because COVID effects tend to be localized by country and sector, renationalization makes firms more dependent on suppliers that are also experiencing the negative shock.

Negative shocks to supply due to lockdowns and global supply chain bottlenecks coincided with a shift in demand away from services and toward goods. Consumption of goods was further boosted by fiscal transfer payments. An increase in household demand coupled with low interest rates resulted in rising inflation rates around the world. While the initial run-up in inflation could be attributed to negative supply shocks, including disruptions in energy markets due to Russia's war on Ukraine, the persistence of inflation suggested that relative price increases were passing through to generalized increases in wages and prices. The challenge facing central banks was to see through the supply-side inflation factors and adjust monetary policy to fight demand-driven inflation. In a sample of 10 countries, Galina Hale, John C. Leer, and Fernanda Nechio find that a 10 percent increase in fiscal support directed to households generated a 40 basis point increase in the inflation rate.[12] Focusing specifically on the United States, Julian di Giovanni, Alvaro Silva, Muhammed A. Yildirim, and Kalemli-Özcan conclude that over the 30 months following the COVID outbreak, shocks to aggregate demand accounted for roughly two-thirds of US inflation, and fiscal stimulus contributed half or more of the demand effect.[13]

Exchange Rate Determination

Exchange rate volatility and the fact that this volatility appears disconnected from other macroeconomic variables have challenged international economists since at least the end of the Bretton Woods era. Recent research finds that nominal exchange rates are primarily driven by shocks originating in financial markets, whether due to noise/liquidity traders, preference shocks, convenience yields, or other shifters. This work has revived a focus that had flourished in the 1970s and 1980s on intermediation frictions. Much of the current work is focused on embedding these frictions in open economy general equilibrium frameworks.

Xavier Gabaix and Matteo Maggiori, for example, construct a general equilibrium model that places financiers at the core.[14] These intermediaries are the only conduit for households in different countries to trade intertemporally or share risk. However, they are subject to a constraint on the size of their net exposure to currency movements, requiring a premium to hold large open positions. This mechanism provides a modern take on the classic portfolio balance model developed by Pentti Kouri. Exchange rates in the model are disconnected from macroeconomic fundamentals and are instead driven by financial shocks to relative currency demand.

In a series of papers, Oleg Itskhoki and Dimitry Mukhin study a general equilibrium framework that nests a range of potential shocks, including shocks to wages, time preference, productivity, government spending, markups, the law of one price, the demand for foreign goods, and international asset demand.[15] The last can proxy for preferences for particular assets, noise trading, risk-bearing capacity, expectation errors, or attitudes toward risk. They argue that international asset demand shocks are crucial in generating exchange rate disconnect. They show this analytically for an economy that approaches an autarky limit as well as computationally for a calibrated version of their general equilibrium model. Their study takes this framework to the collapse of Bretton Woods and shows that it is possible to generate a sharp increase in real exchange rate volatility once a currency begins to float. The key insight is that financial intermediaries are unwilling to bear the risk of completely offsetting noise traders when the exchange rate is volatile and that, conversely, under a peg, intermediaries have a more elastic demand for currencies due to the low-risk environment. Thus, a change in exchange rate regime induces a dramatic shift in the relevance of noise traders. The researchers use their framework to argue that financial market shocks take the lead role in generating a host of exchange rate puzzles, with nominal rigidities — the traditional focus of the literature — recast as a supporting player.

Global Safe Asset/Dollar Pricing

One striking feature about international asset markets is the privileged role of the US dollar. Wenxin Du, Joanne Im, and Jesse Schreger, as well as Zhengyang Jiang, Arvind Krishnamurthy, and Hanno Lustig, show that returns on US Treasuries are consistently below those of similar bonds issued by other G10 economies, even after adjusting for exchange rate risk.[16] In particular, consider the return on US Treasuries minus the yield on a foreign bond plus the difference between the log forward and spot exchange rates in foreign currency per dollar. This is the gap in the covered interest rate parity (CIP) condition for government bonds. If the sign of this expression is negative, then the US bond return is less than a synthetic dollar bond constructed from a foreign bond. Figure 3 plots the time series of the mean and median excess returns on one-year bonds for the US versus nine other developed economies. For comparison, the figure

Figure 3

Deviations from Covered Interest Rate Parity (CIP)

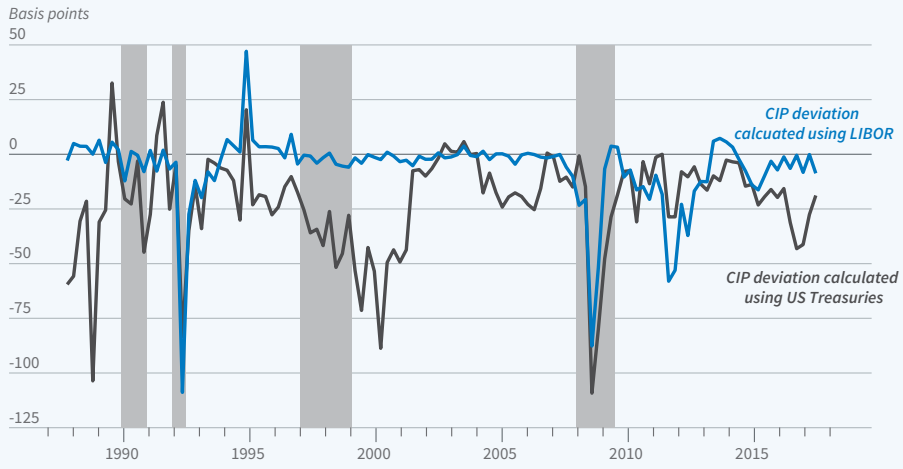

Shaded areas represent recessions as determined by the NBER's Business Cycle Dating Committee.
Source: Zhengyang Jiang, Arvind Krishnamurthy, and Hanno Lustig. NBER Working Paper 24439.

also contains the CIP deviation using the London Interbank Offered Rate (LIBOR), a measure of yields for nongovernment assets. The US Treasury deviation is persistently and significantly negative, with a substantial degree of volatility. Consistent with the important finding of Alexander Tepper, Adrien Verdelhan, and Du, the LIBOR CIP deviation is most prominent after the Global Financial Crisis of 2008.[17] Jiang, Krishnamurthy, and Lustig argue that fluctuations in the implied convenience yield of the dollar are a prime explanatory variable for its exchange rate volatility.

The special role of the dollar contributes to the influence of US monetary policy in the world economy. Silvia Miranda-Agrippino and Hélène Rey document and analyze the fact that capital flows and asset prices around the world co-move strongly with each other.[18] Moreover, a significant driver of these movements is surprises in US monetary policy decisions. Figure 4 shows the impulse responses to a US monetary policy shock for Miranda-Agrippino and Rey's measure of the global factor in asset prices (left panel) and capital flows (right panel). Surprise US monetary policy tightening generates a global decline in asset prices and a retrenchment in international capital flows.

Figure 4

Effect of Tightening US Monetary Policy

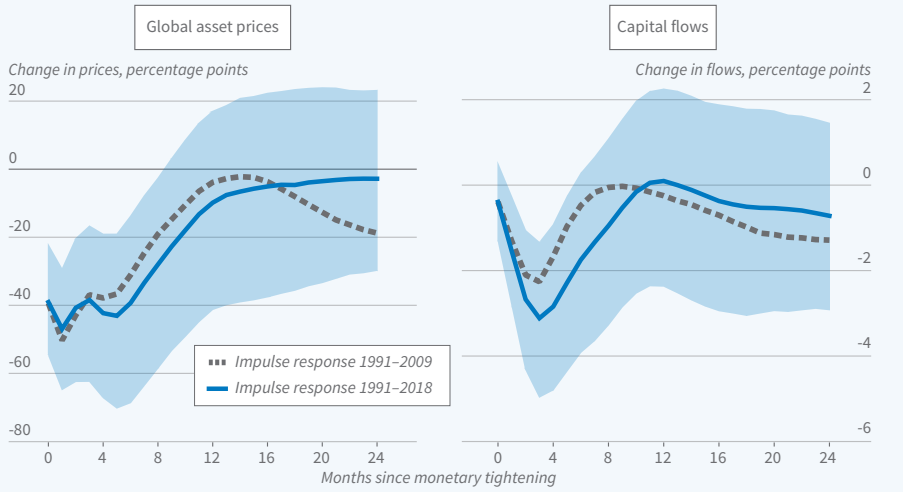

Blue shaded areas represent 90% confidence intervals.
Source: Silvia Miranda-Agrippino and Hélène Rey. NBER Working Paper 29327. Reproduced with permission from Elsevier.

Kalemli-Özcan shows that US monetary policy shocks affect emerging markets and advanced economies differently.[19] In particular, emerging markets tend to match a US rate increase more than one-for-one, while the pass-through for advanced countries is smaller. Despite the increase in the rate differential between emerging markets and the US, capital tends to flow out of emerging markets. Kalemli-Özcan uses additional measures of risk tolerance to argue that the changing appetite for emerging market risk, and the subsequent increase in the equilibrium risk premium, drives the negative correlation between interest rate differentials and capital flows.

1 "A Macroeconomic Model of Healthcare Saturation, Inequality, and the Output-Pandemia Trade-off," Mendoza EG, Rojas EI, Tesar LL, Zhang J. NBER Working Paper 28247, February 2021, and *IMF Economic Review* 71, November 2022, pp. 243–299. Return to Text

2 "COVID-19 Infection Externalities: Trading Off Lives vs. Livelihoods," Bethune ZA, Korinek A. NBER Working Paper 27009, April 2020."Pandemics Through the Lens of Occupations," Chopra A, Devereux MB, Lahiri A. NBER Working Paper 27841, September 2020. "COVID-19 Time-Varying Reproduction Numbers Worldwide: An Empirical Analysis of Mandatory and Voluntary Social Distancing," Chudik A, Pesaran MH, Rebucci A. NBER Working Paper 28629, April 2021.| Return to Text

3 "A Macroeconomic Model of Healthcare Saturation, Inequality, and the Output-Pandemia Trade-off," Mendoza EG, Rojas EI, Tesar LL, Zhang J. NBER Working Paper 28247, February 2021, and *IMF Economic Review* 71, November 2022, pp. 243–299. Return to Text

4 "Deadly Debt Crises: COVID-19 in Emerging Markets," Arellano C, Bai Y, Mihalache GP. NBER Working Paper

27275, March 2023.
Return to Text

5 "EMEs and COVID-19: Shutting Down in a World of Informal and Tiny Firms," Alfaro L, Becerra O, Eslava M. NBER Working Paper 27360, June 2020.
Return to Text

6 "Inside the Mind of a Stock Market Crash," Giglio S, Maggiori M, Stroebel J, Utkus S. NBER Working Paper 27272, May 2020. "Macroeconomic Forecasting Using Filtered Signals from a Stock Market Cross Section," Chatelais N, Stalla-Bourdillon A, Chinn MD. NBER Working Paper 30305, July 2022. "Aggregate and Firm-Level Stock Returns during Pandemics, in Real Time," Alfaro L, Chari A, Greenland AN, Schott PK. NBER Working Paper 26950, May 2020. "External Balance Sheets and the COVID-19 Crisis," Hale G, Juvenal L. NBER Working Paper 29277, September 2021.
Return to Text

7 "The Fed's International Dollar Liquidity Facilities: New Evidence on Effects," Goldberg LS, Ravazzolo F. NBER Working Paper 29982, April 2022. "Monetary Policy Cooperation/Coordination and Global Financial Crises in Historical Perspective," Bordo MD. NBER Working Paper 27898, October 2020.
Return to Text

8 "The Political Economy of the COVID-19 Fiscal Stimulus Packages of 2020," Aizenman J, Jinjarak Y, Nguyen H, Noy I. NBER Working Paper 29360, October 2021.
Return to Text

9 "Policy Packages and Policy Space: Lessons from COVID-19," Bergant K, Forbes K. NBER Working Paper 31254, May 2023.
Return to Text

10 "Fiscal Policy in the Age of COVID: Does It 'Get in All of the Cracks?'," Gourinchas P-O, Kalemli-Özcan Ş, Penciakova V, Sander N. NBER Working Paper 29293, September 2021.
Return to Text

11 "Global Supply Chains in the Pandemic," Bonadio B, Huo Z, Levchenko AA, Pandalai-Nayar N. NBER Working Paper 27224, April 2021.
Return to Text

12 "Inflationary Effects of Fiscal Support to Households and Firms," Hale G, Leer JC, Nechio F. NBER Working Paper 30906, January 2023.
Return to Text

13 "Quantifying the Inflationary Impact of Fiscal Stimulus under Supply Constraints," di Giovanni J, Kalemli-Özcan Ş, Silva A, Yildirim MA. NBER Working Paper 30892, March 2023.
Return to Text

14 "International Liquidity and Exchange Rate Dynamics," Gabaix X, Maggiori M. NBER Working Paper 19854, July 2015.
Return to Text

15 "Exchange Rate Disconnect in General Equilibrium," Itskhoki O, Mukhin D. NBER Working Paper 23401, May 2017. "Mussa Puzzle Redux," Itskhoki O, Mukhin D. NBER Working Paper 28950, June 2021."Optimal Exchange Rate Policy," Itskhoki O, Mukhin D. NBER Working Paper 31933, December 2023. "What Drives the Exchange Rate?" Itskhoki O, Mukhin D. NBER Working Paper 32008, December 2023. This work builds on "Money, Interest Rates, and Exchange Rates with Endogenously Segmented Asset Markets," Alvarez F, Atkeson A, Kehoe PJ. NBER Working Paper 7871, September 2000. "Noise Trading and Exchange Rate Regimes," Jeanne O, Rose AK. NBER Working Paper 7104, April 1999. "Exchange Rates, Interest Rates, and the Risk Premium," Engel C. NBER Working Paper 21042, March 2015.
Return to Text

16 "The US Treasury Premium," Du W, Im J, Schreger J. NBER Working Paper 23759, August 2017. "Foreign Safe Asset Demand and the Dollar Exchange Rate," Jiang Z, Krishnamurthy A, Lustig H. NBER Working Paper 24439, September 2020.
Return to Text

17 "Deviations from Covered Interest Rate Parity," Du W, Tepper A, Verdelhan A. NBER Working Paper 23170, February 2017.
Return to Text

18 "The Global Financial Cycle," Miranda-Agrippino S, Rey H. NBER Working Paper 29327, October 2021. "US Monetary Policy and the Global Financial Cycle," Miranda-Agrippino S, Rey H. NBER Working Paper 21722, March 2019.
Return to Text

19 "US Monetary Policy and International Risk Spillovers," Kalemli-Özcan Ş. NBER Working Paper 26297, September 2019.
Return to Text

Monopsony Power in Labor Markets

Suresh Naidu and Arindrajit Dube

In the standard labor market model taught in introductory economics classes around the world, relationships between firms and workers are just another transaction, mediated by the impersonal market forces of supply and demand. From this perspective, the labor market is best described as being perfectly competitive, where wages are set by the market, with little room for any employer choice.

This "Econ 101" view of the market has a lot of explanatory power. But, as generations of "institutionalist" labor economists have pointed out, it is incomplete. It presumes that the "law of one price" — which postulates that identical commodities have the same price everywhere — holds for workers and their wages. As a result, the model emphasizes the supply of human capital and technology-induced demand as the primary levers that move wages and employment, with little role for firms, power, norms, or interventionist institutions. The competitive model generally portrays efforts to shape the terms of voluntary contracts between workers and firms — like union contracts, minimum wages, or even social conventions — as of secondary importance at best or counterproductive at worst. Fraught arguments about the allocation, and even the definition, of power are unnecessary, as in the words of the late Milton Friedman: "The employee is protected from being coerced by his employer by the existence of other employers for whom he can work."[1]

What are the empirical implications of the theory, and do they match up with the real world? For example, in the hypothetical case of a company reducing wages by 10 percent, the competitive model predicts that all workers eventually will quit and go to competitors. Firms have no scope for wage setting, and the market determines a worker's value at every firm. But is this borne out in the data?

A growing body of experimental and quasi-experimental estimates from a wide variety of contexts suggest otherwise. Even in thick urban labor markets in high-income countries, the share of workers who are likely to leave in response to a hypothetical 10 percent wage cut is much smaller, perhaps 20 to 30 percent, and is often lower for women. In developing economies, it is lower still. This suggests that employers have wide latitude to set wages. A higher wage helps recruit and retain workers, but the market does not dramatically constrain com-

Arindrajit Dube

Arindrajit Dube is a professor of economics at the University of Massachusetts Amherst. He received his PhD in economics from the University of Chicago and his BA in economics and MA in international development policy from Stanford University. Dube's research focuses on the economics of the labor market, including the role of imperfect competition, institutions, norms, and behavioral factors that affect wage setting, minimum wage policies, and jobs. He is an NBER research associate affiliated with the Labor Studies Program. He has testified on the minimum wage before the US Senate Committee on Health, Education, Labor & Pensions, and conducted a review of the evidence on the impact of minimum wages for the government and finance ministry of the United Kingdom, which informed the setting of the level of the National Living Wage.

Suresh Naidu

Suresh Naidu is the Jack Wang and Echo Ren Professor of Economics and a professor of international and public affairs at Columbia University. He received his PhD from the University of California, Berkeley, his master's from the University of Massachusetts Amherst, and a BA in mathematics from the University of Waterloo. Prior to joining Columbia in 2010, he was a Harvard Academy fellow and an instructor at UC Berkeley.

Naidu is an economic historian who focuses on political economy and historical labor markets. His research interests include the economic effects of political transitions, the history of slavery and labor institutions, monopsony, and international migration. He is a member of the external faculty at the Santa Fe Institute, a faculty codirector of the Columbia Center for Political Economy, and an NBER research associate affiliated with the programs in the Development of the American Economy, Political Economy, Development Economics, and Labor Studies.

He has written about housing wealth and political economy. Beyond academic writing, he has contributed to publications such as the *Boston Review*, *Jacobin*, and *The Hindu*.

panies' wage decisions and different employers can make different choices. That is, employers exercise monopsony power — the labor market analog of demand-side monopoly power that gives sellers a degree of control over pricing.

This essay summarizes our recent research on the prevalence of monopsony power and the sources and changing aspects of that power. We also describe our investigation of how monopsony power can affect the provision of nonwage amenities as well as explain a variety of quirks in the labor market, like the widespread use of round-number wage offers or the mixed evidence on the effects of the minimum wage. We end with a discussion of policy and institutional implications of monopsony power as well as new opportunities to study these issues by analyzing the growing set of policy experiments that rebuild countervailing power.

Why Monopsony?

The growing interest in monopsony power fundamentally comes from the wide gap between the theoretical predictions of the competitive model and empirical findings. Perhaps no other prediction is as clearly at odds with the evidence as the "law of one price" in the labor market. Early twentieth-century labor economists marshaled some evidence on this point, but the advent of high-quality, matched employer-employee data, combined with attention to transparent quasi experiments, has moved the research frontier considerably. At this point, there is a large body of evidence for firm-specific wages where workers' wages partly reflect the productivity and profitability of their employers, which violates the law of one price.[2] This naturally pushes us to examine models of the labor market in which employers have some choice in setting wages as opposed to passively taking wages as mere prices of skills, net of disamenities, as dictated by the market.

Some support for monopsony as an influence on wage setting comes from recent studies of the effect of mergers and concentration on wages, which complement earlier evidence on the small effects of the minimum wage on employment. While these indirect failures of the perfect competition model are compelling, the most direct evidence quantifying monopsony measures how firm-level labor supply reacts to plausibly exogenous change in wages. Experimental and quasi-experimental changes in wages across workers, holding all else equal, have been shown to lead to only moderate changes in quits and recruits, a point we will explore in detail in the next section. All this evidence points towards pervasive monopsony as a force in the labor market, where firms set wages for groups of workers, losing those who have better outside options but making profits off those who do not.

What is behind the prevalence of monopsony power in the labor market? Here it is useful to consider the triumvirate of monopsony power: concentration, search frictions, and job differentiation. Concentration refers to having a small number of employers in the market — in the extreme case, a company town. Of course, company towns are rare, but recent research suggests a perhaps surprising amount of concentration in many local labor markets.

However, if concentration were the only source of market power, monopsony power would not be endemic. A key second factor is that job searches are difficult. Large firms set wages for many jobs without having to compete with themselves, creating a bubble of noncompetitive behavior within a workplace. Employees find searching for jobs costly, and many vacancies and possible matches are only communicated informally through social networks.

The third major reason that labor market monopsony exists is that different workers may value the same jobs, paying the same wage, differently. Jobs are more than sources of income and employees factor in more than compensation when making job choices. Many aspects of a job matter, including relationships with coworkers and supervisors, commute times, tastes/abilities for particular tasks, scheduling, and hours. Dube, Naidu, and Adam Reich emphasize that subjective experiences of work — such as meaning and a sense of purpose, managerial respect, and the experience of dignity — are important to workers.[3]

The taste for the same job and knowledge about outside jobs vary for different workers, which gives employers some scope to reduce wages, losing some workers who would rather work elsewhere but keeping those for whom the job is the best they could hope to find.

Measuring Monopsony Power

A central contribution of recent work has been innovative ways of measuring the magnitude of monopsony power. In a paper with Ihsaan Bassier, Dube, and Naidu, we focused on the responsiveness of employee quits in a firm to the wages offered by that firm.[4] In a perfectly competitive labor market, a small increase in wages offered by one employer should attract workers and induce quitting at other firms. With monopsony power, however, firms have some leeway to set wages below competitive levels, as some workers might be reluctant to quit in search of other jobs due to frictions like search and moving costs. Therefore, as Alan Manning points out, the extent of monopsony can be inferred by observing whether quits fall sharply below or only slightly below what would be expected in a competitive model when a firm offers lower wages.[5]

Using administrative employer-employee data from Oregon for information on hourly wages, we devised a matched event-study method in which we consider what happens when two otherwise similar workers initially working at the same firm leave and take jobs at other companies that pay somewhat different wages. Using this approach, illustrated in Figure 1, we estimate the responsiveness of future quits to the firms' wage policies. Typically, if wages are around 10 percent higher, quits are 20–30 percent lower. Generally, this low sensitivity reveals that employers have flexibility in setting wages without fearing excessive turnover. This finding is remarkably consistent across industries, occupa-

Figure 1

Worker Outcomes following a Job Transition

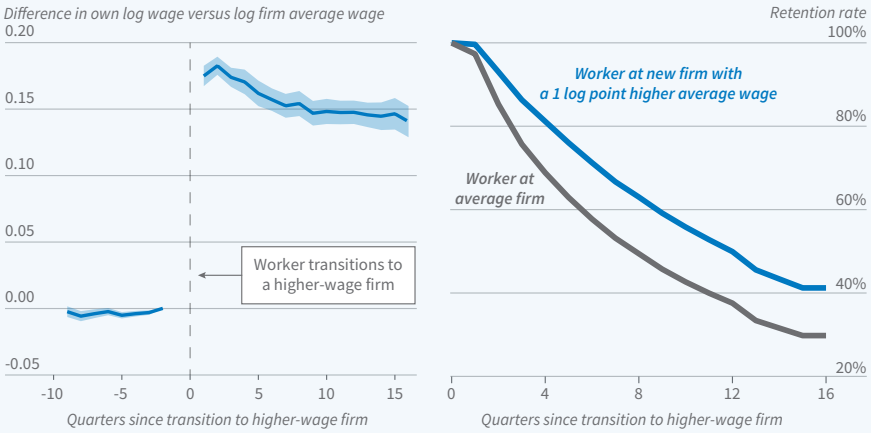

Shaded areas represent 95% confidence intervals.
Source: Ihsaan Bassier, Arindrajit Dube, and Suresh Naidu. NBER Working Paper 27755 and published as "Monopsony in Movers," Journal of Human Resources, 57(S), April 2022, pp. S50–S86. Reproduced under Creative Commons License from The Journal of Human Resources: https://creativecommons.org/licenses/by-nc-nd/4.0/

tions, and geographical regions. Importantly, we find that the quit elasticity is smaller in magnitude for low-wage workers, a finding that has a bearing on designing policies like those for minimum wages.

The Evolution of Monopsony Power

Monopsony power is not static and is affected by macroeconomic conditions. Work by David Autor, Annie McGrew, and Dube highlights dramatic changes in labor market power during the post-pandemic period. An exceptionally tight low-wage labor market following the pandemic significantly enhanced competition.[6] This is reflected in a sharp rise in the quit sensitivity to offered wages, particularly for workers without a college degree, as shown in Figure 2. As a result, workers left lower-paying employers to take better-paying positions, leading to a reinvigorated job ladder. It also likely facilitated the reallocation of workers from lower-productivity to higher-productivity firms.

The heightened responsiveness of quits reduced wage inequality, as low-wage workers saw disproportionate gains. There has been remarkable compression in the 90-10 wage ratio since 2019. The changes of the past four years reversed roughly 40 percent of the increase in the 90-10 ratio between 1980 and 2019. Importantly, much of the compression took place via workers changing their jobs, not from changes in wages at current jobs, highlighting the role of heightened competition.

These findings underscore the importance of overall labor market tightness and macroeconomic policies in determining monopsony power and competition in the labor market, as well as workers' leverage relative to their employers.

Monopsony and Nonwage Amenities

Beyond its immediate impact on wages, monopsony power also has important consequences for the provision of nonwage amenities by employers. Our study with Reich presents a theoretical framework for elucidating these implications.[7] Under perfect competition, competitive forces in the product and labor markets shape optimal decisions for both wages and amenities. However, under monopsony, firms face a trade-off. For instance, if higher wages and the provision of a desired amenity are complements — workers value the amenity more when wages are high — then monopsonistic firms will systematically underprovide that amenity relative to the firms in a competitive market.

We test this idea empirically through stated preference experiments with employees at Walmart. We find that workers value various amenities, including some "dignity-based" amenities relating to fair treatment and predictable scheduling, and find that workers have a stronger appetite for such amenities when wages are higher, lending some support to the complementarity hypothesis.

Next, using variation within our data sample in the extent to which Walmart's $11 per hour corporate min-

Figure 2

Firm Separation and Wage Premium

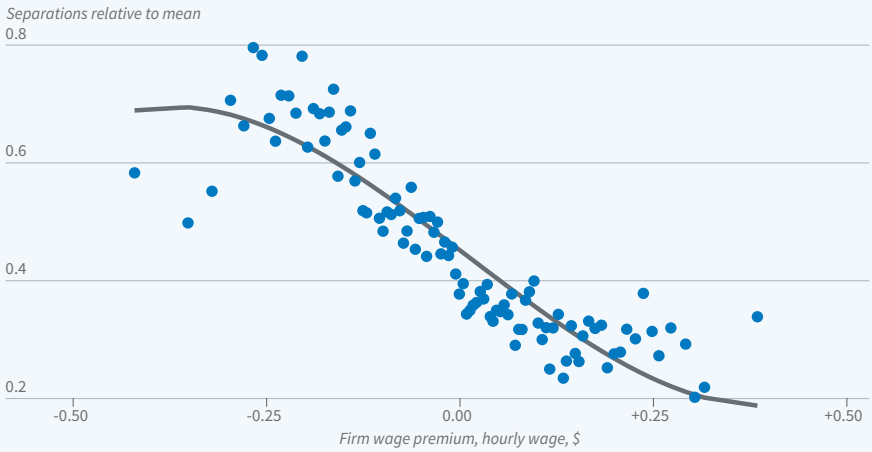

Source: Ihsaan Bassier, Arindrajit Dube, and Suresh Naidu. NBER Working Paper 27755 and published as "Monopsony in Movers," Journal of Human Resources, 57(S), April 2022, pp. S50–S86. Reproduced under Creative Commons License from The Journal of Human Resources: https://creativecommons.org/licenses/by-nc-nd/4.0/

Figure 3

Job-to-Job Transitions for US Non-College Graduates under 40

Source: David Autor, Arindrajit Dube, and Annie McGrew. NBER Working Paper 31010.

imum wage raised wages, we find that more binding minimum wage policies do not lead to a reduction in amenity provision, as might be expected if amenities acted as substitutes for wages. This aligns with the theoretical prediction that firms with monopsony power tend to underprovide amenities that are complementary with wages, especially softer amenities like workplace dignity.

Monopsony, Round-Number Wages, and Mispricing

Intriguingly, the presence of monopsony power can help explain unusual patterns observed in the labor market. The observation that many hourly wages cluster tightly around round numbers, like $10 or $15, has puzzled economists for years. While some explanations center around workers' preference for whole numbers — left-digit bias — our recent research suggests a deeper connection to monopsony power and employer decision-making.

In work coauthored with Alan Manning,[8] we explored this connection, employing a three-pronged approach. The first step involved isolating the source of "round-number bunching." In an online experiment using Amazon Mechanical Turk, we offered slightly different, nonround wages for identical tasks. If workers had a strong left-digit bias, there should have been a sharp increase in worker acceptance at the round-number wage compared to the just-below-round wage. However, we found no such discontinuity, suggesting workers did not treat round numbers differently. At the same time, we uncovered considerable monopsony power using experimental evidence in this online labor market, as acceptance rates were only modestly related to the offered wage.

The lack of worker left-digit bias points to the possibility that the bunching may reflect employer preferences or mispricing. To assess this, we leveraged high-quality matched employer-employee data. If employers were strategically targeting round numbers due to worker preference, one would expect most of the "excess jobs" at the round-numbered wage — say, $10 — to be coming from below $10. However, empirically we found that the "missing jobs" were coming from both above and below $10, contradicting the left-digit bias explanation but consistent with employer mispricing. This showed that less-sophisticated employers were more likely to use round-numbered wages.

What is behind such mispricing? Some employers, particularly less sophisticated ones, might face optimization frictions, so round numbers might be a convenient, low-cost heuristic for setting wages. Critically, we show theoretically and empirically that such mispricing is much more likely to exist in markets with some monopsony power. In a competitive market, small deviations from the optimal wage result in significant profit losses, pushing firms towards precise wage setting. However, with monopsony power, the relationship between profits and wages becomes flatter around the optimal wage level. This flatter envelope means that even moderate deviations from the optimal wage might not significantly harm profits, allowing some firms to settle for the ease and famil-

Figure 4

Probability of Quitting Given an External Offer

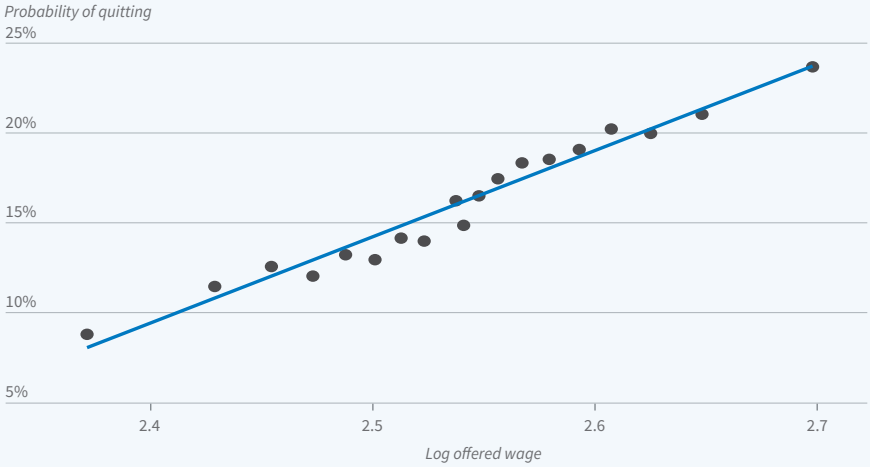

Movement from 2.5 to 2.6 represents approximately a 10 percent increase in the offered wage.
Arindrajit Dube, Suresh Naidu, and Adam D. Reich. NBER Working Paper 30441.

Figure 5

Hourly Wages in Minnesota, Oregon, and Washington, 2003–07

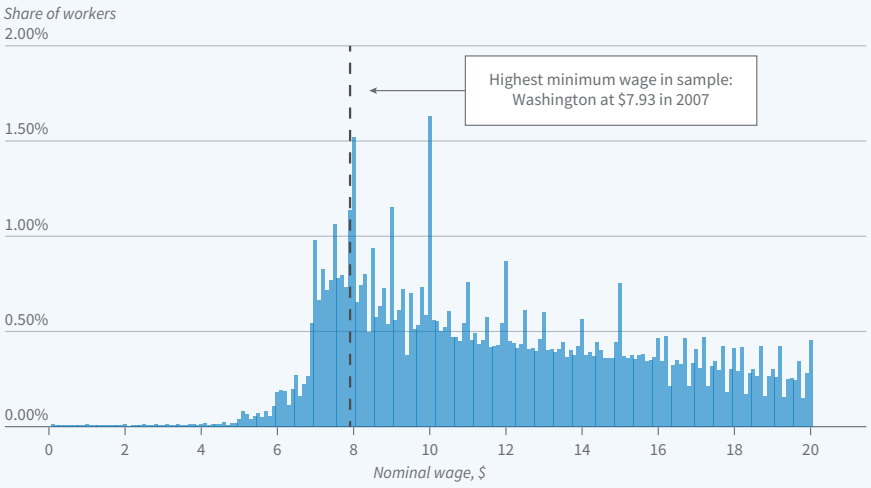

Source: Arindrajit Dube, Alan Manning, and Suresh Naidu. NBER Working Paper 24991.

iarity of round numbers.

This research suggests that, beyond round numbers, clunky wage policies and quirky pay rules are easier to explain in labor markets with pervasive monopsony.

The Rebirth of Countervailing Power?

If monopsony is indeed pervasive, policies that offer countervailing power to employer-side monopsony can increase wages and improve the functioning of the economy. Policies that might provide countervailing power include antitrust policies to reduce concentration, macro policies to increase labor market tightness and quit elasticity, and minimum wages and collective bargaining to reduce employer wage-setting power. Whether these mechanisms work remains an empirical question. The good news for researchers seeking to answer this question is that policymakers are experimenting along many of these dimensions.

When minimum wages are imposed or raised, the number of low-productivity jobs may shrink, but labor market monopsony implies that high-productivity jobs will expand, and so the overall effect on employment will be small. This is consistent with recent evidence from work by Doruk Cengiz, Attila Lindner, Ben Zipperer, and Dube.[9]

An inherent limitation of a minimum wage policy is that it only affects the bottom of the wage distribution. Collective and sectoral wage bargaining between employers and democratic unions have the potential to improve efficiency, fairness, and the balance of power in the labor market more broadly. In the US, there is growing state-level experimentation with standards in the fast-food and healthcare sectors, for example, in California and Minnesota. These can offer important insights into the efficacy of sectoral standards in improving worker welfare in monopsonistic markets.

More directly, as discussed by Naidu, Eric Posner, and Glen Weyl, we are seeing increased use of antitrust tools to counter labor market concentration and explicitly anticompetitive practices.[10] The most recent Department of Justice merger guidelines explicitly mention harms from labor market power, and the Federal Trade Commission has recently proposed a ban on noncompete clauses in employment contracts. We are also seeing private class-action lawsuits charging employers with unfair competition in the labor market, for example, the ongoing $1.6 billion suit filed by 1,200 fighters against Ultimate Fighting Championship.

Besides particular policies focused on wages and labor markets, underlying monopsony power moves with the broader macroeconomic environment. The role of tightness in reducing monopsony power raises the possibility that macro policy can move labor market power at a much broader level, and integrating this insight into the design of optimal monetary and fiscal policy may prove important.

These recent policy innovations, at scales ranging from firms to markets to whole economies, probe the feasibility of various modes of countervailing power. They complement a recent in-

Figure 6

Employee Separation Rates in Oregon, by Wage Bin

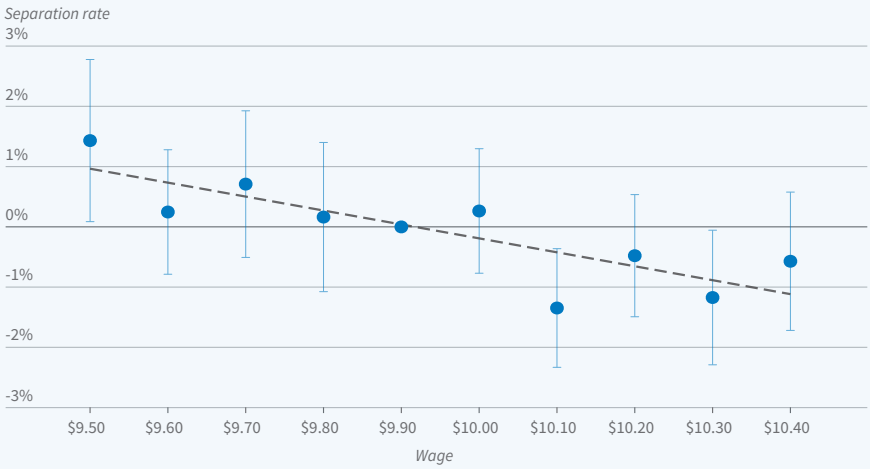

Bars represent 95% confidence intervals.
Source: Arindrajit Dube, Alan Manning, and Suresh Naidu. NBER Working Paper 24991.

terest in the political economy of "predistribution" policies that intervene in the labor market, as contrasted with the traditional focus of economists on tax-and-transfer policies.[11] These novel policies also provide economists with a host of naturally occurring experiments for learning about the forces shaping employer wage setting.

Sections of this summary draw substantially on a March 2024 article in the International Monetary Fund's Finance & Development magazine, authored by Suresh Naidu.

1 *Capitalism and Freedom*, Friedman M. Chicago: University of Chicago Press, 1962.
Return to Text

2 "Monopsony in the Labor Market: New Empirical Results and New Public Policies," Ashenfelter OC, Card D, Farber HS, Ransom M. NBER Working Paper 29522, November 2021.
Return to Text

3 "Power and Dignity in the Low-Wage Labor Market: Theory and Evidence from Walmart Workers," Dube A, Naidu S, Reich AD. NBER Working Paper 30441, September 2022.
Return to Text

4 "Monopsony in Movers: The Elasticity of Labor Supply to Firm Wage Policies," Bassier I, Dube A, Naidu S. NBER Working Paper 27755, August 2020.
Return to Text

5 "The Real Thin Theory: Monopsony in Modern Labour Markets," Manning A. *Labour Economics* 10(2), April 2003, pp. 105–131.
Return to Text

6 "The Unexpected Compression: Competition at Work in the Low Wage Labor Market," Autor D, Dube A, McGrew A. NBER Working Paper 31010, November 2023.
Return to Text

7 "Power and Dignity in the Low-Wage Labor Market: Theory and Evidence from Walmart Workers," Dube A, Naidu S, Reich AD. NBER Working Paper 30441, September 2022.
Return to Text

8 "Monopsony and Employer Mis-optimization Explain Why Wages Bunch at Round Numbers," Dube A, Manning A, Naidu S. NBER Working Paper 24991, January 2020.
Return to Text

9 "The Effect of Minimum Wages on Low-Wage Jobs: Evidence from the United States Using a Bunching Estimator," Cengiz D, Dube A, Lindner A, Zipperer B. NBER Working Paper 25434, January 2019.
Return to Text

10 "Antitrust Remedies for Labor Market Power," Posner EA, Naidu S, Weyl G. *Harvard Law Review* 132(2), December 2018, pp. 536–601.
Return to Text

11 "'Compensate the Losers?' Economic Policy and Partisan Realignment in the US," Kuziemko I, Longuet-Marx N, Naidu S. NBER Working Paper 31794, October 2023.
Return to Text

Economic Incentives in Pay-for-Performance Programs

Edward Norton

The Centers for Medicare and Medicaid Services (CMS) spends nearly $1 trillion per year on healthcare expenditures for Medicare beneficiaries. With such large payments to healthcare providers, CMS is concerned about promoting quality of care. Over the last few decades, it has created several programs that reward hospitals and other providers financially for achieving measurable outcomes.[1,2] These are commonly known as pay-for-performance, or P4P, programs. Their goal is to give providers larger financial payments in the future if current quality measures are high or improving.

The economic issues addressed in P4P programs are challenging. Without any quality incentives, there is concern that providers would strive to increase the quantity of care without regard to quality, leading to higher total outlays but not necessarily health improvement. To combat this moral hazard problem, CMS wants to link payments in part to quality of care. The challenge is to make financial incentives large enough to encourage improvements but not so large as to cause other distortions. In addition, if the P4P programs were too punitive to low-performing providers, the programs could lead to hospital closures, potentially lowering access to care in already underserved areas.

My research is about the economic incentives in Medicare's P4P programs. In particular, I am interested in measuring financial incentives at the patient level to see if the distribution of these incentives is related to hospital-level characteristics, and in discovering whether changes in quality over time are related to the incentives. The fundamental assumptions of P4P programs are that providers have financial incentives to improve care and that they respond to those incentives. My research tests those assumptions.

Moneyball in Medicare

My interest in P4P programs in healthcare goes beyond studying whether these programs change quality of care or spending after implementation. Such an analysis could be done with a difference-in-differences empirical analysis. Instead, I am primarily concerned with understanding the economic incentives in the programs and whether they align with the goals. Most of my research on this topic has focused on Medicare's Hospital Value-Based Purchasing (Hospital VBP) program.

The Hospital VBP program measures outcomes in four broad domains: patient experience, clinical outcomes, mortality and safety, and episode spending, defined as any healthcare spending between the admission that launches the episode and 30 days post-discharge. The program is important because nearly 3,000 general hospitals participate in it, and up to 3 percent of a hospital's future Medicare reimbursement depends on its performance each year.

Jun Li, Anup Das, Lena Chen, and I tested the two fundamental assumptions of P4P for the Hospital VBP program.[3] First, how large are the financial incentives? This turns out to be hard to estimate because the incentives vary by domain and across patients as well as across hospitals. Heterogeneity is an important feature of these incentives. We established that Medicare payment for one patient hospitalization is not just the diagnosis-related group (DRG) payment, as it was prior to Hospital VBP, but now includes a marginal future reimbursement equal to how that patient's outcomes affect the hospital's VBP score, rating, and future payment, where treatment occurred. Simply put, a hospital's total Medicare reimbursement for one patient is the sum of its current DRG payment and the discounted marginal future reimbursement based on how that patient's outcomes affect the future bonus.

For example, take the mortality domain, which measures 30-day mortality for acute myocardial infarction, heart failure, and pneumonia. A patient's outcome affects the mortality rate, which affects the number of points received for the mortality measure, which affects the total score across all domains, which affects the bonus percentage, which affects future Medicare reimbursement for the hospital. The hospital's future reimbursement is affected by whether the patient lives or dies.

Edward C. Norton

Edward C. Norton is the UnitedHealthcare Professor of Health Care Management at the University of Michigan and a professor in both the School of Public Health and the Department of Economics. He is an NBER research associate affiliated with the Economics of Health and Economics of Aging Programs. Norton's research interests include health economics, applied econometrics, long-term care and aging, and pay-for-performance.

Norton received his AB from Princeton University and his PhD from MIT. He has held positions at Harvard Medical School, RTI International, and the University of North Carolina at Chapel Hill, and visiting positions at the London School of Economics, University of Tokyo, Hunter College, and Dartmouth College. He has received funding for his research from several government agencies.

When we calculated the change in future annual Medicare payments at the hospital level due to a hypothetical patient death, we found that the marginal effect on future reimbursement was not always negative, as we expected it would be. Instead, we found that it was zero for about one-third of hospitals. Due to the complex nonlinear incentives, a sizable fraction of hospitals faced no penalty for worsening mortality. Similarly, they received no financial benefit if they improved mortality slightly. For hospitals with nonzero incentives the median financial benefit for avoiding a patient death was less than $10,000. For a few hospitals the value was larger, sometimes as much as $40,000.

The heterogeneity in mortality incentives was similar to what we found in other domains. Improving the quality of current patient outcomes had no effect on marginal future reimbursements for the hospitals that treated between a quarter and a third of patients. P4P in practice has a wide range of financial incentives across hospitals, with a sizable fraction facing no meaningful financial incentive to improve quality of care at the margin.

We also tested the second P4P assumption, which is that hospitals respond to the incentives. There were several reasons to be skeptical of hospital responses, beginning with the fact that there is a lag of about two years between quality measurements and the application of bonuses or penalties. Also, clinical personnel making treatment decisions do not directly receive any financial rewards, and professional norms promote quality even without financial incentives. Finally, the amount of the bonus or penalty might be too small to affect behavior. Despite such concerns, the entire premise of P4P is that the way to achieve better quality of care is to pay hospitals to improve.

Our evidence supports the presence of some behavioral response.[4] We tested whether the year-over-year change in each quality measure was related to the marginal future reimbursement — technically, the marginal change in the total performance score given a one decile change in that measure, a measure of the magnitude of the incentive. Of the 15 measures we tested, seven were statistically significant and of the expected sign.

Our framework can be used to analyze any of the P4P programs, not just Hospital VBP. It remains an open question whether the same results would be found in, for example, the Hospital Readmissions Reduction Program or the Hospital-Acquired Condition Reduction Program.

Heterogeneous Treatment Effects

While our original research established a relationship between marginal financial incentives and year-over-year improvement in measures, the exact nature of the relationship was unclear. Emily Lawton, Li, and I next turned our attention to measuring the functional form of the relationship between these two variables.[5] This can reveal the cost-effectiveness of the Hospital VBP program. To visualize this relationship, imagine graphing the change in quality of care on the *y*-axis as a function of the marginal future reimbursement on the *x*-axis.

If the true relationship is along a straight line from the origin, then each hospital has the same ratio of marginal financial incentive to change in quality of care. Small incentives lead to small improvements in quality of care, and large incentives lead to proportionately larger improvements. To obtain a high return on its investment, CMS wants hospitals to be in the upper right corner of the graph, that is, to make large improvements in quality of care for a small financial incentive. In contrast, CMS wants to avoid paying large incentives and getting little or no change in quality in return, as represented by the lower left part of Figure 1.

Another possibility for the relationship could be a discontinuous jump at the origin, with small financial incentives discontinuously inducing modest increases in quality, and then perhaps a concave function for positive values. Finally, there can be no relationship at all if the program is too confusing or hospitals are focused on other issues. It could be that hospitals ignore the incentives and if by random luck they happen to improve measured quality of care anyway, then they are happy to collect their bonus payment. In that case, the hospitals would be scattered along the *x*-axis with no apparent relationship.

Empirically we found that larger financial incentives induce better outcomes. However, this relationship is not linear. There is a large jump at zero when the incentives become positive. The large discontinuous jump implies

Figure 1

Links between Change in Hospital Quality and Future Reimbursement

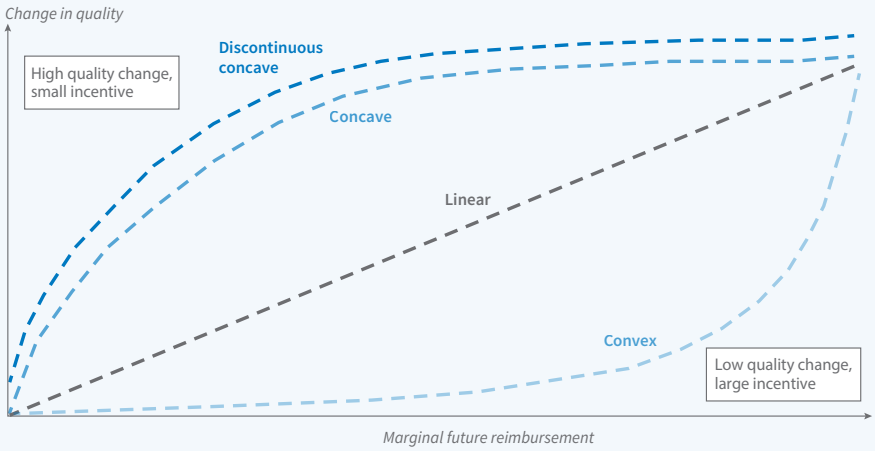

Source: Edward C. Norton, Emily J. Lawton, and Jun Li. "Moneyball in Medicare: Heterogeneous Treatment Effects," NBER Working Paper 27948 and in American Journal of Health Economics *9(1), Winter 2023, pp. 96–126. Reproduced with permission from the* American Journal of Health Economics.

Figure 2

Tradeoff between Mortality Reduction and Cost

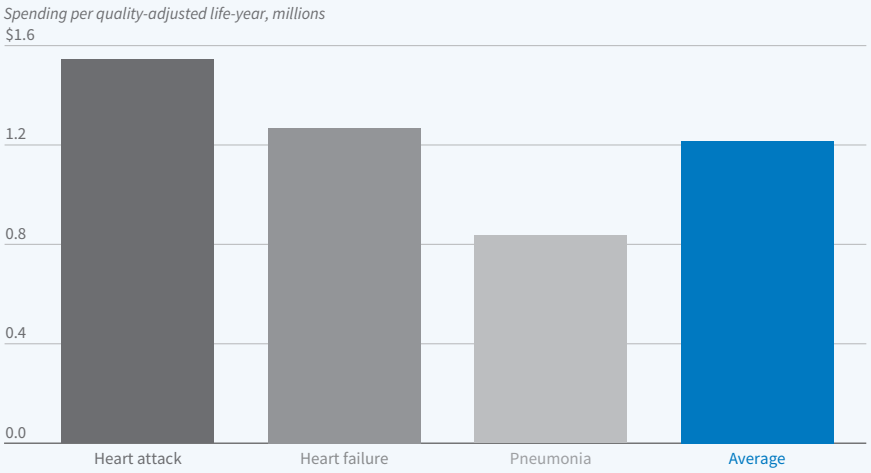

Source: Edward C. Norton, Jun Li, Anup Das, Andrew M. Ryan, Lena M. Chen. "Medicare's Hospital Value-Based Purchasing Program Values Quality over QALYs," Medical Decision Making *42(1), January 2022, pp. 51–59.*

that small positive incentives can induce hospitals to improve quality of care. This is the more cost-effective side of the figure.

Hospital VBP Values Quality over QALYs

Unlike P4P programs that target a single outcome, the Hospital VBP program measures a broad range of outcomes, from patient satisfaction and clinical outcomes to mortality and spending. Li, Das, Andy Ryan, Chen, and I explored the economic ramifications of the trade-offs implicit in such a composite measure.[6] The Hospital VBP program converts scores on many quality measures into points and ultimately into dollars. The formulas that make that conversion are production functions. Patient outcomes create points. A hospital can earn more points either by doing better on, say, the patient heart attack mortality rate or on total episode spending.

We were interested in estimating the magnitude of the trade-off between improved mortality and lower spending. An optimizing hospital has a choice, at the margin, of lowering the mortality rate or spending less to earn the same number of points. When seen as a production function, the Hospital VBP program implicitly trades off between lives and dollars, and we calculated that trade-off. What is the improvement in mortality necessary to earn points, and what is the corresponding improvement in spending needed to earn the same number of points?

Quality of care is often measured in quality-adjusted life-years (QALYs), where a QALY is a measure of the quality and quantity of life lived, with 1.0 QALY being one full year of life in perfect health. The commonly accepted range for medical interventions is roughly $50,000 to $200,000 per QALY. If incentives in the Hospital VBP program are balanced, then the trade-off between spending improvement and mortality improvement measured in QALYs should be in this range.

We estimate the total value of Medicare savings divided by the equivalent total of QALYs gained. These findings imply that the value of the mortality reduction in the Hospital VBP program is $1,542,837 per QALY for heart attack, $1,268,827 per QALY for heart failure, and $835,129 per QALY for pneumonia. The average across all three conditions is $1,215,598 per QALY. These numbers are several orders of magnitude higher than the accepted range, which suggests that the Hospital VBP program overvalues improvements in quality of care, relative to spending reductions, relative to what we judge to be the common accepted valuation metrics.

1 "Incentive Regulation of Nursing Homes," Norton EC. *Journal of Health Economics* 11(2), August 1992, pp. 105–128.
Return to Text

2 "Do Report Cards Predict Future Quality? The Case of Skilled Nursing Facilities," Cornell PY, Grabowski DC, Norton EC, Rahman M. NBER Working Paper 25940, June 2019, and *Journal of Health Economics* 66, May 2019, pp. 208–221.
Return to Text

3 "Moneyball in Medicare," Norton EC, Li J, Das A, Chen LM. NBER Working Paper 22371, June 2016, and *Journal of Health Economics* 61, September 2018, pp. 259–273.
Return to Text

4 Ibid.
Return to Text

5 "Moneyball in Medicare: Heterogeneous Treatment Effects," Norton EC, Lawton EJ, Li J. NBER Working Paper 27948, October 2020, and *American Journal of Health Economics* 9(1), Winter 2023, pp. 96–126.
Return to Text

6 "Medicare's Hospital Value-Based Purchasing Program Values Quality over QALYs," Norton EC, Li J, Das A, Ryan AM, Chen LM. *Medical Decision Making* 42(1), January 2022, pp. 51–59.
Return to Text

The Economics of Generative AI

Erik Brynjolfsson and Danielle Li

Artificial intelligence (AI) is not a new field. The term was coined in 1956, but the field has only recently begun having significant effects on the economy.

Research in AI went through three eras. Early work focused primarily on symbolic systems with hand-coded rules and instructions. In the 1980s, expert systems, which consisted of hundreds or thousands of "if...then" rules drawn from interviews with human experts, helped diagnose diseases and make loan recommendations, but with limited commercial success.

Later, the focus shifted to machine learning systems, including "supervised learning" systems trained to make predictions based on large datasets of human-labeled examples. As computational power increased, deep learning algorithms became increasingly successful, leading to an explosion of interest in AI in the 2010s.

More recently, even larger models using unsupervised or self-supervised systems have become a major focus of the field. Large-language models (LLMs) — trained on massive amounts of text to simply predict the next word in a sequence — have astounded the public with their ability to produce meaningful and remarkable output. These systems have been found to outperform humans for a growing range of knowledge-intensive tasks: taking the bar exam, for instance. In addition, studies show that access to LLMs and other types of generative AI tools can help human workers improve their own performance.

In the past year, a growing body of work has explored how new AI tools might impact productivity in applications as diverse as coding, writing, and management consulting.[1]

In research with Lindsey Raymond, we analyze the effects of generative AI on worker productivity in the context of technical customer support.[2] Our study is based on data from over 5,179 agents, about 1,300 of whom were given access to an LLM-based assistant that provided real-time suggestions for communicating with customers. The system, trained on millions of examples of successful and unsuccessful conversations, provided suggestions that the agents could use, adapt, or reject. The tool was rolled out in phases, creating quasi-experimental evidence on its causal effects.

We found significant improvements in worker productivity as measured by the number of customer issues workers were able to resolve per hour. Within

Erik Brynjolfsson

Erik Brynjolfsson is the Jerry Yang and Akiko Yamazaki Professor and Senior Fellow at the Stanford Institute for Human-Centered Artificial Intelligence, and director of the Stanford Digital Economy Lab. He also is the Ralph Landau Senior Fellow at the Stanford Institute for Economic Policy Research, professor by courtesy at the Stanford Graduate School of Business and Stanford Department of Economics, and an NBER research associate affiliated with the Productivity, Innovation, and Entrepreneurship Program. He was previously a professor at MIT Sloan.

Brynjolfsson's research focuses on the economics of AI and other digital technologies. He has measured the productivity contributions of information technology and the complementary role of intangible assets. Brynjolfsson holds bachelor's and master's degrees from Harvard University in applied mathematics and decision sciences and a PhD from MIT in managerial economics.

Danielle Li

Danielle Li is an associate professor in the Technological Innovation, Entrepreneurship, and Strategic Management group at the MIT Sloan School of Management, and an NBER faculty research fellow affiliated with the program in Productivity, Innovation, and Entrepreneurship. Her work focuses on how organizations consume information to make decisions: how do firms decide what projects to invest in or which people to hire and promote, and what are the economic consequences of those decisions? Her past work has examined the impact of taking advice from biased experts on the quality of R&D funding, as well as the impact of algorithm design on hiring outcomes in the labor market.

Li has previously taught at Harvard Business School and the Kellogg School of Management at Northwestern University. She holds a PhD in economics from MIT and an AB in mathematics and the history of science from Harvard College.

Figure 1

Productivity of Customer Support Agents and AI Support

Bars represent 95% confidence intervals.
Source: Erik Brynjolfsson, Danielle Li, and Lindsey R. Raymond. NBER Working Paper 31161.

four months, treated agents were outperforming nontreated agents who had been on the job for over twice as long.

On average, worker productivity increased by 14 percent. These gains were concentrated among the lowest quintile of workers, whether measured by experience or prior productivity, where there were productivity gains of up to 35 percent. In contrast, the top quintile saw negligible gains and, in some cases, even small decreases in the quality of conversations, as measured by customer satisfaction. This pattern is reflective of how the system is trained: by observing successful conversations, the system is able to glean the behavior of the most skilled agents and pass on these behaviors as suggestions to novice workers.

Did the system deskill the workforce? Another natural experiment suggests not. As with most large systems, there were occasional outages when the system unexpectedly became unavailable. Workers who had previously been using the system now had to answer questions without access to it, and nonetheless they continued to outperform those who had never used the system. This suggests that the system helped them learn, and retain, answers.

Our results point to the possibility that — in contrast with earlier waves of information technology that largely benefited higher-skill workers — generative AI technologies could particularly benefit workers at the lower or middle levels of the skills distribution. Drawing on these and other results, David Autor sees opportunities for the recent waves of AI to help rebuild the middle class by increasing the value of output from their labor.[3]

Advances in AI technologies and algorithmic design can yield improvements beyond direct measures of productivity. For example, we saw evidence in our study that AI assistance improves the experience of work for treated agents, as measured by the processing of conversation transcripts: customers spoke more kindly to agents and were less likely to ask to speak to a supervisor. These effects were likely driven both by agents' improved social skills and increased access to technical knowledge as a result of chat assistance.

Indeed, there is growing evidence that generative AI tools may outperform humans in an area traditionally considered a source of strength for humans relative to machines: empathy and social skills. One study of doctors' responses to patient questions found that an LLM-based chatbot provided answers that were judged by expert human evaluators to be more detailed, higher quality, and 10 times more likely to be considered empathetic.[4]

Finally, innovations in AI systems may further improve the functioning of current AI tools. For example, Li, Raymond, and Peter Bergman explore how algorithm design can improve the quality of interview decisions in the context of professional services hiring. They find that while traditional supervised learning systems — which look for workers who match historical patterns of success in the firm's training data — select higher-quality workers relative to human hiring, they are also far less likely to select applicants who are Black or Hispanic. In contrast, reinforcement learning and contextual bandit models — which value learning about workers who have not traditionally been represented in the firm's training data — are able to deliver similar improvements in worker quality while also distributing job opportunities more broadly.

While the effects of AI on productivity and work practices are now evident not only in a number of laboratory settings but also in business applications, it may take longer for them to show up in aggregate statistics. Brynjolfsson, Daniel Rock, and Chad Syverson discuss a set of reasons why the effects of AI might not quickly change aggregate productivity numbers.[5]

For one thing, labor productivity is typically defined as GDP per hour worked. But GDP as it is traditionally measured may miss many of the benefits of an increasingly digital economy that creates free goods and makes them more widely available while also improving the quality, variety, or convenience of existing goods. An alternative metric, GDP-B, seeks to address these challenges by assessing the benefits of goods and services, not the amount spent.[6]

Furthermore, general purpose technologies like AI are likely to experience a lag between their initial adoption and observable improvements in productivity. In a second study, Brynjolfsson, Rock, and Syverson model this "Productivity J-Curve."[7] As with other types of information technology, the initial phase of AI adoption is characterized by time-consuming complementary investments, including the realignment of business processes, the integra-

Figure 2

Productivity Mismeasurement J-Curve

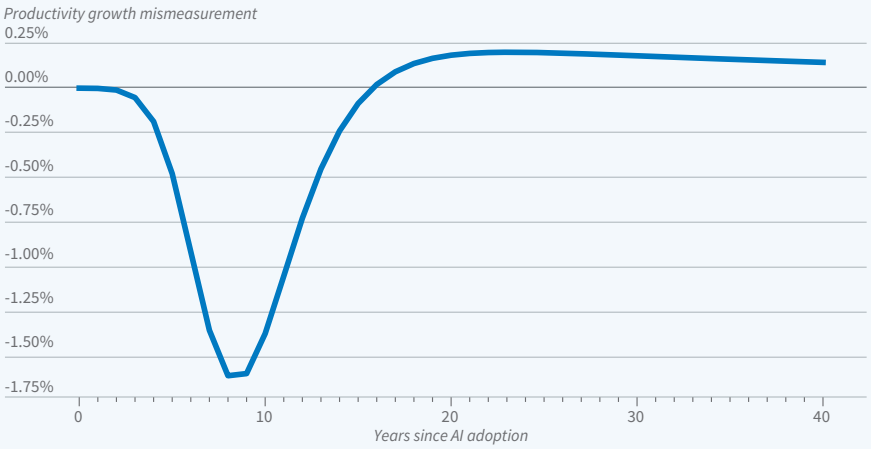

Source: Erik Brynjolfsson, Daniel Rock, and Chad Syverson. NBER Working Paper 25148, and published as "The Productivity J-Curve: How Intangibles Complement General Purpose Technologies," American Economic Journal: Macroeconomics 13(1), January 2021, pp. 333–72.

tion of new technologies into existing workflows, and the upskilling of the workforce. As noted by Brynjolfsson and Lorin Hitt, these adjustments are costly and may create valuable intangible assets, but neither the costs nor the benefits are typically accounted for when measuring a firm's output.[8] As a result, productivity as it is conventionally measured may initially be seen as stagnating or even falling. However, as these technological and organizational complements are gradually implemented, the productivity benefits of AI begin to materialize, marked by an upward trajectory in the J-curve.

The Productivity J-Curve model implies that productivity metrics fail to capture the full extent of benefits during the initial stages of AI adoption, leading to underestimation of AI's potential.

The ultimate economic effects of generative AI will depend not only upon how much it boosts productivity and changes work in specific cases, but also on how much of the economy it is likely to affect. As noted by Daron Acemoglu and Autor, occupations can be broken down into specific tasks.[9] Applying this insight, Brynjolfsson, Tom Mitchell, and Rock look at 18,156 tasks in the O-NET taxonomy and find that most occupations include at least some tasks that could be automated or augmented by machine learning, though significant redesign would typically be required to realize the full potential of the technology.[10] Building on this work, Tyna Eloundou, Sam Manning, Pamela Mishkin, and Rock estimate that approximately 80 percent of the US workforce could have at least 10 percent of their work tasks either automated or augmented by the introduction of LLMs, while around 19 percent of workers could see at least half of their tasks affected.[11]

Hulten's theorem states that a first-order approximation of the productivity effects of a technology is the share of the economy affected multiplied by its average productivity impact. There is evidence that both the potential productivity impact and the potential share of the economy affected are significant in the case of generative AI, suggesting that the ultimate effects may be substantial, though, as implied by the Productivity J-Curve, they may take some time to be realized.[12]

The field of economics itself is not immune to the effects of generative AI. Students of economics are using the tools to help with their assignments, requiring a rethinking of teaching methods. We and our colleagues are using the tools to help with research and writing; we used LLMs to help with aspects of the preparation of this article. Anton Korinek described six ways that LLMs can assist economists: ideation and feedback, writing, background research, data analysis, coding, and mathematical derivations.[13] Jens Ludwig and Sendhil Mullainathan go further, showing that AI models can be used to make the first stage of the scientific process — hypothesis generation — more systematic.[14]

As discussed by Brynjolfsson and Gabriel Unger, important policy choices are emerging regarding AI's effects on productivity, industrial concentration, and inequality.[15] For instance, on the question of inequality, the distinction between technology used for automation versus augmentation or, more formally, AI that substitutes for rather than complements labor, can have significant effects on the distribution of income and bargaining power.[16] Brynjolfsson has argued that either approach can boost productivity but has noted that a focus on human-like AI can lead to a "Turing Trap" by reducing worker bargaining power. As AI continues to grow in power, so too does the need for economic research to better understand how we can harness its benefits while mitigating its risks.

1 "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence," Noy S, Zhang W. *Science* 381(6654), July 2023, pp. 187–192. "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality," Dell'Acqua F, McFowland III E, Mollick E, Lifshitz-Assaf H, Kellogg KC, Rajendran S, Krayer L, Candelon F, Lakhani KR. Harvard Business School Working Paper No. 24-013, September 2023.
Return to Text

2 "Generative AI at Work," Brynjolfsson E, Li D, Raymond LR. NBER Working Paper 31161, November 2023.
Return to Text

3 "Applying AI to Rebuild Middle Class Jobs," Autor D. NBER Working Paper 32140. February 2024.
Return to Text

4 "Comparing Physician and Artificial Intelligence Chatbot Responses to Patient Questions Posted to a Pub-

lic Social Media Forum," Ayers JW, Poliak A, Dredze M, Leas EC, Zhu Z, Kelley JB, Faix DJ, Goodman AM, et. al. *JAMA Internal Medicine* 183(6), April 2023, pp. 589–596.
Return to Text

5 "Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics," Brynjolfsson E, Rock D, Syverson C. NBER Working Paper 24001, November 2017.
Return to Text

6 "GDP-B: Accounting for the Value of New and Free Goods in the Digital Economy," Brynjolfsson E, Collis A, Diewert WE, Eggers F, Fox KJ. NBER Working Paper 25695, March 2019.
Return to Text

7 "The Productivity J-Curve: How Intangibles Complement General Purpose Technologies," Brynjolfsson E, Rock D, Syverson C. NBER Working Paper 25148, January 2020, and *American Economic Journal: Macroeconomics* 13(1), January 2021, pp. 333–372.
Return to Text

8 "Beyond Computation: Information Technology, Organizational Transformation and Business Performance," Brynjolfsson E, Hitt LM. *Journal of Economic Perspectives*, 14(4), Fall 2000, pp. 23–48.
Return to Text

9 "Skills, Tasks and Technologies: Implications for Employment and Earnings," Acemoglu D, Autor D. NBER Working Paper 16082, June 2010. Published as "Chapter 12 - Skills, Tasks and Technologies: Implications for Employment and Earnings" in *Handbook of Labor Economics* 4(B), 2011, pp. 1043–1171.
Return to Text

10 "What Can Machines Learn, and What Does It Mean for Occupations and the Economy?" Brynjolfsson E, Mitchell T, Rock D. *AEA Papers and Proceedings* 108, May 2018, pp. 43–47.
Return to Text

11 "GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models," Eloundou T, Manning S, Mishkin P, Rock D. *arXiv*, August 2023.
Return to Text

12 "Machines of Mind: The Case for an AI-Powered Productivity Boom," Baily MN, Brynjolfsson E, Korinek A. Brookings Institution, May 10, 2023.
Return to Text

13 "Generative AI for Economic Research: Use Cases and Implications for Economists," Korinek A, *Journal of Economic Literature* 61(4), December 2023, pp. 1281–1317.
Return to Text

14 "Machine Learning as a Tool for Hypothesis Generation," Ludwig J, Mullainathan S. NBER Working Paper 31017, March 2023.
Return to Text

15 "The Macroeconomics of Artificial Intelligence," Brynjolfsson E, Unger G. International Monetary Fund, December 2023.
Return to Text

16 "The Turing Trap: The Promise & Peril of Human-Like Artificial Intelligence," Brynjolfsson E. *Daedalus* 151(2), Spring 2022, pp. 272–287. An earlier version of this argument was published as *Race Against the Machine: How the Digital Revolution is Accelerating Innovation, Driving Productivity, and Irreversibly Transforming Employment and the Economy*, Brynjolfsson E, McAfee A. Digital Frontier Press, 2011.
Return to Text

Unemployment in Informal Labor Markets in Developing Countries

Emily Breza and Supreet Kaur

Developing countries typically exhibit low rates of rural wage employment. For example, in India, male workers whose primary source of earnings is wage labor report working on only 46 percent of days per year.[1] Bangladesh has a similarly low 55 percent rate of employment among landless males, and the rates are even lower in sub-Saharan Africa.[2]

What do these low employment rates mean? One possibility is that they reflect extremely high involuntary unemployment. Alternatively, the rates could be an outcome of reasonably well functioning labor markets in which workers are simply choosing self-employment, which tends to be high in poor countries. These two possibilities have drastically different implications for understanding how well labor markets work and what role, if any, there is for policy intervention.

Our work has sought to characterize the functioning of labor markets in developing countries and examine microfoundations for why markets might not always be clear. In this summary, we focus on rural labor markets, with evidence primarily drawn from India. These markets are of intrinsic interest: they are the primary source of wage employment for over a billion people worldwide, including the world's poorest, and their features — informal, decentralized spot markets for labor, where workers and employers match for short-term contracts — are ubiquitous in both rural and urban areas of poor countries. Consequently, many, though not all, lessons from this work likely apply more broadly in developing-country settings.

Staggering Involuntary Unemployment

We begin with the central question of whether low employment rates reflect involuntary or voluntary unemployment. We tackle this question with Yogita Shamdasani by developing a new empirical approach.[3] We induce transitory hiring shocks in randomly selected Indian local labor markets — villages — by "removing" on average 24 percent of male workers by giving them factory jobs in sites outside of the village for two to four weeks. This

Emily Breza

Emily Breza is the Frederic E. Abbe Professor of Economics at Harvard University. She received her PhD from MIT and her BA from Yale University. Prior to Harvard, she held a faculty appointment at Columbia Business School.

Breza is a development economist whose research focuses on social networks and their consequences for labor markets and financial markets. She is a board member and finance cochair of the Abdul Latif Jameel Poverty Action Lab (J-PAL) and an NBER research associate affiliated with the programs in Development Economics; Productivity, Innovation, and Entrepreneurship; and Corporate Finance. She has faculty affiliations at the International Growth Centre and the Centre for Economic Policy Research (CEPR).

Breza is the recipient of a Sloan Research Fellowship and a National Science Foundation CAREER award.

Supreet Kaur

Supreet Kaur is an associate professor in the Department of Economics at the University of California, Berkeley. She is a development economist whose work applies insights from psychology to the study of poverty and labor markets.

Kaur is a recipient of a Sloan Research Fellowship, a National Science Foundation CAREER award, and the Distinguished CESifo Affiliate Award in Behavioral Economics. She was awarded the David A. Wells Prize in 2012 by the Harvard Economics Department for best thesis embodying the results of original investigation. She holds a PhD in political economy and government (Economics) from Harvard, a master's in international development from the Harvard Kennedy School, and a BS in operations research from Columbia University. She codirects the Psychology and Economics of Poverty Initiative at Berkeley's Center for Effective Global Action. She is an NBER research associate affiliated with the Development Economics and Labor Studies Programs.

shock substantively reduces the number of workers in the local village economy without changing labor demand within the village. Looking at the local labor market response to our hiring shock enables us to learn how many people wanted a job at the prevailing wage but could not find one before we intervened. Importantly, we learn this simply by observing the employment behavior of the remaining workers and employers.

We find distinctly different patterns of effects in "lean" versus "peak" months, reflecting seasonality in agricultural hiring. In lean months we detect severe rationing: at least one in four workers in the economy wants a wage job but cannot find one. Excluding our external factory jobs, removing a quarter of the labor force has no effect on lean season wages or aggregate employment. This is consistent with rationed workers filling the newly vacated job slots, leaving the total number of people holding a job unchanged. In contrast, in peak months the impact of our hiring shock matches a textbook supply and demand model: wages rise quickly — within a week — and aggregate employment in the village falls, so that each new job created in the external factory jobsites generates only 0.74 days of new work for laborers in the economy overall. Together, these findings present a more nuanced picture of informal labor markets: they have the capacity to be extremely agile and responsive, but also exhibit extremely high levels of labor rationing in lean months.

Our approach contrasts with how economists have typically measured unemployment to date: asking people in surveys whether they were looking for work but could not find it. It is unclear whether such survey self-reports are reliable.[4] By basing our measurement on whether workers actually choose to work when job slots in the local economy open up, we obtain revealed preference estimates of rationing.

Our approach also lets us learn about self-employment. We find that many rationed workers are disguised as entrepreneurs: as soon as job slots open up in their village, entrepreneurs readily abandon their agricultural and nonagricultural businesses in order to take a wage job with a local employer. In lean months, at least 24 percent of self-employment stems from workers being rationed out of wage labor. Among farmers with small landholdings, this shift to self-employment is especially high: in lean months, at least 64 percent of work on small farms would not occur if those farmers could find wage jobs instead. Consequently, our results indicate that a substantial fraction of self-employment stems from poor labor market prospects rather than high growth opportunities. This can help us understand why broadly targeted interventions such as credit, wage subsidies, and training for microenterprises tend to generate low average returns.

These patterns indicate why answers to standard involuntary unemployment questions can be unreliable in developing countries, and more broadly in settings with self-employment or informal work like gig jobs. In employment surveys run by governments — in India, the US, and most other settings — workers are only classified as involuntarily unemployed if they are not engaged in any work activity. Since a rationed worker who cannot find wage work can turn to self-employment or a gig economy job to make some extra money, focusing on these standard questions can lead to drastic underestimation of labor rationing in the economy. We show that alternate employment status questions that take this into account can yield more accurate estimates.

But why is there so much labor rationing? If there are more workers who want jobs than there are jobs available, we would expect wages to fall until the supply of workers equals demand. Kaur documents that in rural India, wages are downwardly rigid.[5] Specifically, while they rise in response to positive shocks, they do not fall in response to negative shocks, such as droughts. Kaur's study shows that downward rigidity causes increased unemployment — arguably the first direct evidence of employment effects of wage rigidity in any setting.

In addition to wage rigidity, wage compression also seems to be exhibited in labor markets in this setting: workers of varying abilities are paid the same wage. As we show in Figure 1, there tends to be a single prevailing wage for a given type of task in the economy, which most workers are paid despite differences in underlying ability. One consequence is that all workers agree on what the prevailing wage is for a task, a feature that plays an important role in labor market dynamics.

Figure 1

Daily Wage Distribution in Rural Orissa, India

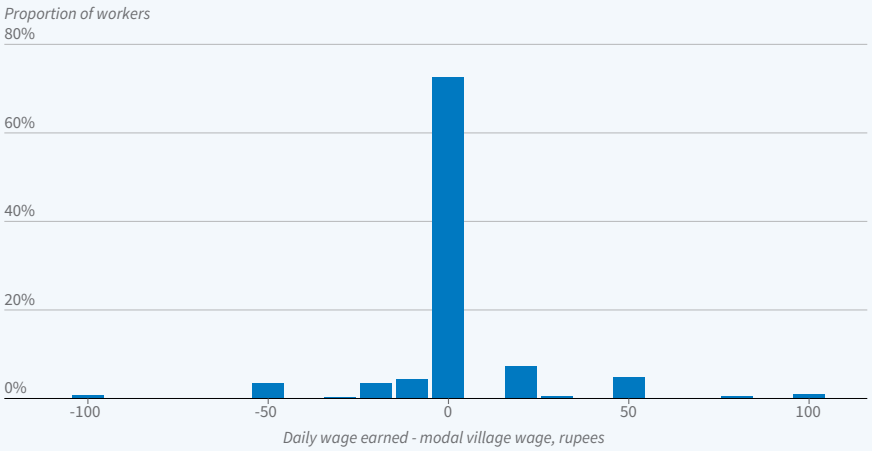

Source: Emily Breza, Supreet Kaur, and Yogita Shamdasani. NBER Working Paper 22491, and published as "The Morale Effects of Pay Inequality," The Quarterly Journal of Economics, *133(2), May 2018, pp. 611–663. Reproduced with the permission of* The Quarterly Journal of Economics.

Market Power, Unemployment, and Social Forces

Informal labor markets often embody features we might associate with competitiveness and flexibility: there are many decentralized workers and employers, the short duration of contracts means that wages could quickly reflect changes in market conditions, there are no formal unions or institutions, and minimum wage laws are generally ignored. Why, then, do wages seem inflexible — over time, across people, and in response to shocks? Understanding this is key for understanding the high levels of unemployment.

Our research indicates that it is essential to account for one additional feature of these markets: workers are not anonymous to one another.[6] They live in close-knit communities and are dependent on each other socially and economically — for example, through job referrals and informal insurance. This creates scope for individuals to use the threat of sanctions to sustain norms that are perceived to be in their collective interest.

In work with Nandita Krishnaswamy, we document norms against accepting wage or price cuts in a range of markets in India and Kenya.[7] In Figure 2, we show evidence from rural agricultural workers, construction workers at urban labor stands, taxi drivers, food vendors, and butchers. In each case, respondents state that undercutting the prevailing price is considered unacceptable [Panel A]. Doing so would result in a range of sanctions, from being socially ostracized to losing one's source of livelihood [Panel B]. For example, 90 percent of rural workers respond that others would get angry with an individual who accepted a job below the prevailing wage and 59 percent believe others would impede the labor market prospects of that individual by means such as withholding referrals.

We then use a field experiment to examine whether norms against accepting wage cuts can help us understand the presence of wage floors in rural labor markets. Specifically, we hypothesize that during times of unemployment, at least some workers would prefer working below the prevailing wage rather than remaining jobless, but do not due to the threat of sanctions from other workers.

We partner with local employers who make job offers for lean season agricultural work. Employers follow the typical process for agricultural hiring: traveling to the neighborhood where the majority of workers live and making job offers to workers at their homes. We randomize both the wage level of the offer and the degree of observability.

In the "public" treatment, which is the status quo for our setting, the employer offers the job outside on the street where neighbors, who are typically other workers, can overhear the offer, and could then tell other workers in the community. In the "employer" treatment, the wage offer is observable to the employer but not to others in the community. In the "private" condition, the job is offered inside the worker's home and consequently is not directly observable to others in

Figure 2a

Social Acceptability of Wage and Price Cuts

Survey results on acceptability of cutting product prices or wages by 10 percent.

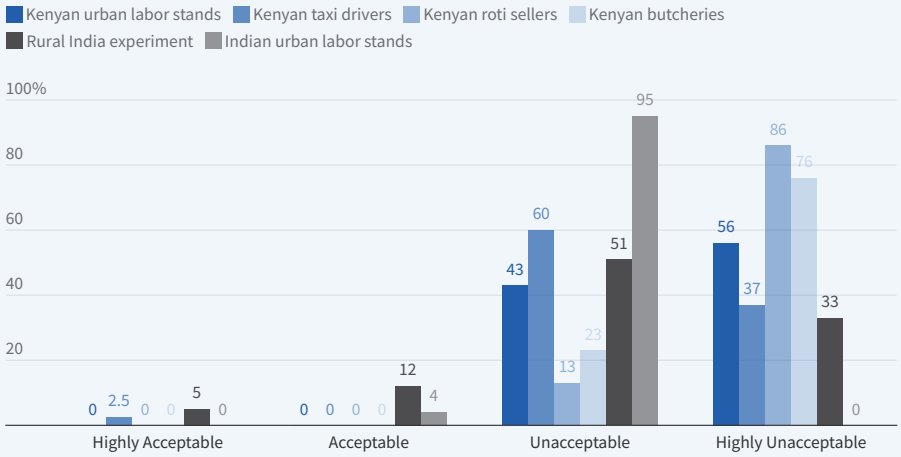

Labor stands are places where day laborers gather to be hired.

Source: Emily Breza, Supreet Kaur, and Nandita Krishnaswamy. NBER Working Paper 25880.

Figure 2b

Community Responses to Wage and Price Cutters

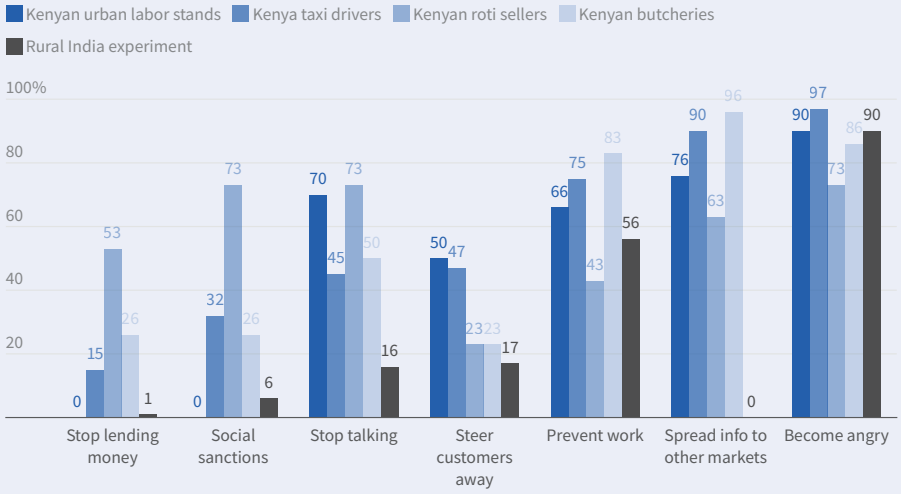

Respondents were read a list of options and could select any that apply.

Source: Emily Breza, Supreet Kaur, and Nandita Krishnaswamy. NBER Working Paper 25880.

Figure 3

Agricultural Labor Supply at Submarket Wages

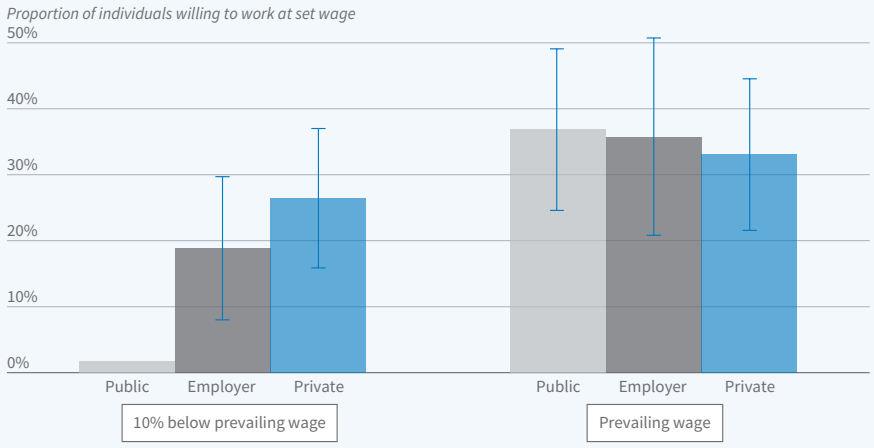

In a public setting, wages are known to all workers. In an employer setting, wages are known only to the employer. In a private setting, wages are known only to the worker. Bars represent 90% confidence intervals. Source: Emily Breza, Supreet Kaur, and Nandita Krishnaswamy. NBER Working Paper 25880.

the community. After the conclusion of the experiment, participants received a supplement so that no one actually worked below the prevailing wage.

Despite high unemployment, only 1.8 percent of agricultural workers are willing to accept jobs below the prevailing wage under the status quo. However, this number jumps to 26 percent when this choice is not observable to other workers. In contrast, for prevailing-wage job offers, social observability has no detectable impact on job take-up. This is consistent with the idea that social observability only matters when there are norm violations.

These findings are consistent with substantial distortions in the aggregate supply curve. At low wages, social pressure leads to a restriction in labor supply, making it appear that below the prevailing wage, labor supply falls to close to zero. However, absent social considerations, unemployed workers would prefer working below the prevailing wage to remaining jobless. Whether the norm among workers increases or decreases, total employment depends on whether employers have market power. Regardless, a back-of-the-envelope calculation suggests that maintaining a wage floor is beneficial for worker earnings as a whole.

Overall, our findings suggest that market power can arise in many decentralized markets and may be more widespread in developing countries than has been previously thought.

Wage Compression: Social Forces

Social forces also have relevance for explaining wage compression. With Shamdasani, we explore whether and under what circumstances workers care about their pay relative to that of their coworkers.[8] If relative pay differences cause workers to withdraw labor supply or effort, employers may prefer offering compressed wage contracts.

We conduct an experiment with workers in seasonal, month-long, low-skilled manufacturing jobs in India. Workers are randomly organized into three-person production units, each of which is randomized to one of four different pay structures. In the pay disparity condition, each unit member is paid a different wage — w_{High}, w_{Med}, or w_{Low} — in accordance with his respective productivity rank within the unit determined by baseline productivity levels. In the three compressed pay conditions, all unit members are paid the same wage, which we randomly assign to be w_{High}, w_{Med}, or w_{Low}. This allows us to compare, for example, workers with the same average baseline productivity who both earn an absolute wage of w_{Low}, but differ in whether they are paid less than their peers under the pay disparity treatment or the same as their peers under the compressed low wage treatment.

Figure 4 shows the impacts of pay disparity on standardized output holding own wage fixed. Prior to "Day 0" of the experiment, all workers were paid identical training wages. For low-ranked workers earning w_{Low}, output declines by 0.33 standard deviations (22 percent) on average after a work-

Figure 4

Pay Disparity and Output

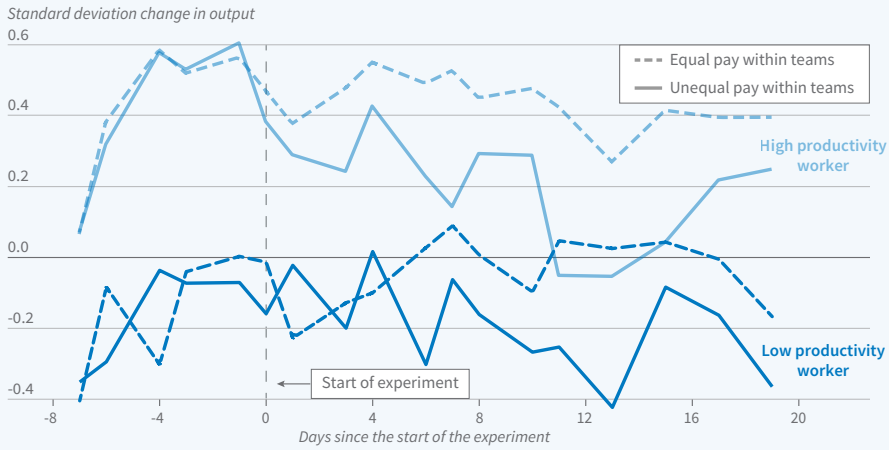

Source: Emily Breza, Supreet Kaur, and Yogita Shamdasani. NBER Working Paper 22491, and published as "The Morale Effects of Pay Inequality," The Quarterly Journal of Economics, 133(2), May 2018, pp. 611–663. Reproduced with the permission of The Quarterly Journal of Economics.

er begins earning less than both his coworkers, and attendance declines by 12 percentage points. Perhaps surprisingly, the high-ranked workers in pay disparity units, who earn *more* than both their coworkers, also experience a reduction in output and labor supply.

We find that perceived justifications play an important role in mediating the effects. When workers can clearly perceive that their higher-paid peers are more productive than themselves, pay disparity has no discernible negative effects on output or labor supply. That workers tolerate pay inequality only when productivity differences are extremely transparent can help explain why we observe piece rates in practice where output is fully observable but do not often observe other forms of pay dispersion.

Finally, we show that these morale effects likely operate through resentment and hostility in the workplace, reducing social cohesion among unit members. In two incentivized, cooperative games, members of pay disparity units are less able to work together than members of compressed units, even when it's in their own interest. However, in both cases, when pay disparity is clearly justified, these effects disappear.

Together, this body of work makes progress toward understanding the functioning of rural labor markets in developing countries. It shows that while these markets embody remarkable flexibility and agility, wages are rigid and involuntary unemployment is extremely high, particularly during months when agricultural labor demand is low. This changes the logic of labor market analysis. For example, because wages do not always play an allocative role, analyses that assume wages tell us something about labor productivity will be misleading. In addition, our findings are relevant for poverty alleviation policies. For example, they suggest that workfare programs that offer jobs to unemployed workers — a popular policy tool in developing countries — are least likely to crowd out private sector jobs if implemented in lean seasons, but may do so in shoulder or peak seasons.

Why such high levels of unemployment exist in this setting remains an open question. Researchers have failed to find support for the traditional microfoundations that were discussed in the early development literature, such as nutrition efficiency wages. Our work makes progress on this puzzle by highlighting a microfoundation not previously considered: the centrality of social forces in determining market outcomes. Because markets are made up of people, they are underpinned by social relationships that can drastically alter their functioning.

1 "India - National Sample Survey 2009–2010," Ministry of Statistics and Programme Implementation. International Labour Organization, January 21, 2024.
Return to Text

2 "Effects of Emigration on Rural Labor Markets," Akram AA, Chowdhury S, Mobarak AM. NBER Working Paper 23929, October 2017. *Accelerating Poverty Reduction in Africa*, Beegle K, Christiaensen L, editors. Washington, DC: The World Bank, October 2019.
Return to Text

3 "Labor Rationing," Breza E, Kaur S, Shamdasani Y. NBER Working Paper 28643, April 2021, and *American Economic Review* 111(10), October 2021, pp. 3184–3224.
Return to Text

4 "Measurement Error in Survey Data," Bound J, Brown C, Mathiowetz N. In *Handbook of Econometrics*, volume 5, Heckman J, Leamer E, editors, pp. 3705–3843. Amsterdam: Elsevier, 2001. "Involuntary Unemployment," Taylor JB. In *The New Palgrave Dictionary of Economics*. London: Palgrave Macmillan, 2008.
Return to Text

5 "Nominal Wage Rigidity in Village Labor Markets," Kaur S. NBER Working Paper 20770, November 2017, and *American Economic Review* 109(10), October 2019, pp. 3585–3616.
Return to Text

6 "The Morale Effects of Pay Inequality," Breza E, Kaur S, Shamdasani Y. NBER Working Paper 22491, August 2016, and *The Quarterly Journal of Economics* 133(2), May 2018, pp. 611–663.
Return to Text

7 "Social Norms as a Determinant of Aggregate Labor Supply," Breza E, Kaur S, Krishnaswamy N. NBER Working Paper 25880, February 2024.
Return to Text

8 "The Morale Effects of Pay Inequality," Breza E, Kaur S, Shamdasani Y. NBER Working Paper 22491, August 2016, and *The Quarterly Journal of Economics* 133(2), May 2018, pp. 611–663.
Return to Text

Annual Report of Awards to NBER Affiliates, Spring 2024

Anjali Adukia received the Advancing Justice, Equity, Diversity, and Inclusion Award from the Association for Education Finance & Policy, and the Society for Research on Educational Effectiveness Early Career Award.

Joseph G. Altonji received the Jacob Mincer Award from the Society of Labor Economists for lifetime contributions to labor economics.

Elizabeth Ananat was appointed to the National Academies of Sciences, Engineering, and Medicine's Standing Committee on Reproductive Health, Equity, and Society.

Joshua Angrist, **Hilary Hoynes**, **Monika Piazzesi**, and **Emmanuel Saez** were elected to the National Academy of Sciences.

Orley Ashenfelter received an honorary doctorate from the University of Bordeaux.

Orazio Attanasio was elected to the American Academy of Arts and Sciences and awarded an honorary doctorate at the Università di Padova, where he delivered a lecture in the hall where Galileo once taught.

David Autor was awarded the NOMIS Distinguished Scientist Award.

Matthew Backus, **Christopher Conlon**, and **Michael Sinkinson** received the *AEJ: Microeconomics* Best Paper Award for "Common Ownership in America: 1980–2017."

Katherine Baicker received the Reinhardt Distinguished Career Award.

Martin Beraja, **Eduardo Dávila**, and **Ellora Derenoncourt** each received National Science Foundation CAREER Awards.

Alan Blinder received the Daniel Patrick Moynihan Prize from the American Academy of Political and Social Science for "demonstrat[ing] the value of using research and evidence to improve the human condition."

David Bloom received the Population Association of America's Irene B. Taeuber Award for "…important contribution to the scientific study of population or for an accumulated record of exceptionally sound and innovative research." Also, with coauthors Maddalena Ferranna, Lisa A. Robinson, Daniel Cadarette, and Michael R. Eber, he received *Risk Analysis'* Best Article Award for "The Benefits and Costs of US Employer COVID-19 Vaccine Mandates."

Markus Brunnermeier received the Ludwig Erhard Prize for Business Journalism.

Guillermo A. Calvo, **Olivia S. Mitchell**, **Maurice Obstfeld**, and **Christina Romer** were named 2023 American Economic Association Distinguished Fellows.

Amitabh Chandra was elected to the American Academy of Arts and Sciences.

John H. Cochrane was awarded a Bradley Prize by the Lynde and Harry Bradley Foundation, and *The Economist* named his book, *The Fiscal Theory of the Price Level*, one of the best books of 2023.

Lauren H. Cohen was awarded a Fulbright Visiting Scholarship to aid Cambodia in building their business education system.

Janet Currie served as president-elect of the American Economic Association, president of the Western Economic Association International, was a Distinguished CES Fellow, and received the Klaus J. Jacobs Research Prize for work aimed at improving learning, development, and living conditions of children.

Manasi Deshpande, **Peter Hull**, **Alex Imas**, **Yueran Ma**, **Tobias Salz**, and **Guo Xu** were each awarded a Sloan Research Fellowship.

Janice Eberly and Nicolas Crouzet received the Dimensional Fund Advisors Best Paper Prize for their paper in *The Journal of Finance*, "Rents and Intangible Capital: A Q+ Framework."

Alex Eble was awarded a Jacobs Foundation Research Fellowship.

Mark Egan, Shan Ge, and Johnny Tang won the TIAA Paul A. Samuelson Award for "Conflicting Interests and the Effect of Fiduciary Duty: Evidence from Variable Annuities." Egan, **Adi Sunderam**, and Stefan Lewellen received the Michael J. Brennan Best Paper Award from the Society for Financial Studies for their paper, "The Cross-Section of Bank Value" in *The Review of Financial Studies*.

Price Fishback was awarded the Gallman-Parker Prize from the Economic History Association for compiling data and information and sharing it with other scholars.

Caroline Flammer was awarded the inaugural University of Oxford Greening Finance Prize for her research on sustainable finance.

Nicole Fortin became a Fellow of the Canadian Economics Association.

Claudia Goldin received the 2023 Sveriges Riksbank Prize in Economic Sciences in Memory of Alfred Nobel and an honorary Doctorate of Science from the University of Rochester.

Mark Grinblatt, Söhnke Bartram, and Yoshio Nozawa won the Best Paper Award from the European Capital Markets Institute for "Book-to-Market, Mispricing, and the Cross-Section of Corporate Bond Returns."

Gordon Hanson was named an International Economic Association Fellow for 2023.

Zhiguo He, **Hui Chen**, Zhuo Chen, Jinyu Liu, and Rengming Xie received a *Journal of Finance*-Dimensional Fund

Advisors Distinguished Paper Prize for "Pledgeability and Asset Prices: Evidence from the Chinese Corporate Bond Markets."

Oliver Hart was appointed Knight Bachelor in the King's Birthday Honours list.

David Hirshleifer was inducted as a Fellow of Financial Management Association International.

Jill Horwitz won the American Bar Association Section of Business Law Committee on Nonprofit Organization's 2023 Outstanding Academic Award.

Daniel Hungerman was named a 2024/25 Phi Beta Kappa Visiting Scholar.

Yunan Ji was awarded the Young Economists' Essay Award at the European Association for Research in Industrial Economics for "Can Competitive Bidding Work in Health Care? Evidence from Medicare Durable Medical Equipment."

Lawrence Jin and Cary Frydman received the Vernon L. Smith Excellence Award in Experimental Finance from the Society for Experimental Finance for "Efficient Coding and Risky Choice."

Louis Kaplow received the American Law and Economics Association's Ronald H. Coase Medal for major contributions to law and economics.

Anil Kashyap was appointed an Honorary Commander of the British Empire for his services to the UK economy.

Amanda Kowalski received the 2023 ASHEcon Willard G. Manning Memorial Award for the Best Research in Health Econometrics.

Steven Lehrer received the 2023 Dan Usher Prize for Excellence in Economic Research.

Derek Lemoine and **Ashley Langer** were awarded the Ralph C. d'Arge and Allen V. Kneese Award for Outstanding Publication in the *Journal of the Association of Environmental and Resource Economists* for "Designing Dynamic Subsidies to Spur Adoption of New Technologies."

Christian Leuz received an honorary doctorate from Maastricht University in the Netherlands.

Nadya Malenko received the Referee of the Year Award from the Society for Financial Studies.

Ulrike Malmendier received an honorary doctorate from Leuphana University Lüneburg, the Award of Excellence from the German American Business Association, and the Skandia Award from the Thule Foundation.

Randall Morck was awarded an honorary doctorate by the University of Montreal/HEC Montréal.

Nitya Pandalai-Nayar received the Kiel Institute Excellence Award in Global Economic Affairs.

Philip Oreopoulos was elected a Fellow of the Society of Labor Economists.

Lindsay Page, Darryl Hill, Rodney Hughes, Matthew Lenard, and David Liebowitz received a best paper award from the *Economics of Education Review* for "New Schools and New Classmates: The Disruption and Peer Group Effects of School Reassignment."

Ariel Pakes received the Erwin Plein Nemmers Prize in Economics from Northwestern University.

Vincent Pons received the Best Young Economist of France Award from *Le Monde* and Le Cercle des économistes.

Alessando Rebucci was awarded a Bank of England Houblon-Norman and George Fellowship.

Maya Rossin-Slater received the Elaine Bennett Research Prize from the AEA Committee on the Status of Women in the Economics Profession.

Philipp Schnabl, Daniel Paravisini, and Veronica Rappoport received the *Journal of Finance* Brattle Group First Prize for "Specialization in Bank Lending: Evidence from Exporting Firms."

Kathryn E. Spier received an honorary doctorate from BI Norwegian Business School.

Stijn Van Nieuwerburgh served as past president of the American Real Estate and Urban Economics Association.

Dimitri Vayanos was elected a Fellow of the Finance Theory Group.

David Weinstein was inducted to the Order of the Rising Sun, Gold Rays with Neck Ribbon, one of the highest honors bestowed by the Japanese government.

Motohiro Yogo was elected a Fellow of the Econometric Society.

Joshua S. Graff Zivin was named a Web of Science Highly Cited Researcher, ranking in the top 1 percent of the citation distribution within economics and business.

Gabriel Zucman was awarded the John Bates Clark Medal by the American Economic Association.

Conferences and Meetings, Spring 2024

Detailed programs for NBER conferences are available at nber.org/conferences

Title of Conference/Meeting	Organizers	Dates
Economics of Education Program Meeting	Caroline M. Hoxby	November 30–December 1, 2023
Big Data and Securities Markets	Itay Goldstein, Chester S. Spatt, and Mao Ye	December 1, 2023
Entrepreneurship Working Group	Josh Lerner and David T. Robinson	December 1, 2023
Development Economics/BREAD Program Meeting	Nava Ashraf, Lorenzo Casaburi, Paulina Oliva, Benjamin A. Olken, Imran Rasul, and Dean Yang	December 1, 2023
International Trade and Investment Program Meeting	Cecile Gaubert and Andrés Rodríguez-Clare	December 1–2, 2023
Innovation Information Initiative Technical Working Group	Adam B. Jaffe	December 1–2, 2023
Compensation of Top Executives: Determinants and Consequences	Dirk Jenter and Kelly Shue	December 7, 2023
Economics of Health	Christopher S. Carpenter, Monica Deza, Amy Finkelstein, and Tal Gross	December 7–8, 2023
Design and Regulation of Transportation Markets	Giulia Brancaccio, Nicholas Buchholz, and Tobias Salz	December 8, 2023
Innovative Data in Household Finance: Opportunities and Challenges	Pascal J. Noel, Jialan Wang, and Stephen P. Zeldes	December 8, 2023
The Economics of Decarbonizing Industrial Production	Lint Barrage and Kenneth Gillingham	December 8–9, 2023
Chinese Economy Working Group Meeting	Hanming Fang, Zhiguo He, Shang-Jin Wei, and Wei Xiong	December 15–16, 2023
Capital Markets, Technology, Financial Inclusion, and Economic Growth	Shilpa Aggarwal and Amit Seru	December 16–17, 2023
Mentorship Program to Support NSF Grant Proposal Development for MSI Faculty Workshop	Danielle Dickens, James M. Poterba, and Angelino Viceisza	January 25, 2024
Economics of Supply Chains	Laura Alfaro and Chad Syverson	January 26, 2024
Longer-Term Health and Economic Effects of COVID–19	Gopi Shah Goda and Maria Polyakova	February 2, 2024
Economics of Crime Working Group Meeting	Crystal Yang and Jens Ludwig	February 2, 2024
Industrial Organization Program Meeting	John Asker, Matthew Backus, Shoshana Vasserman, and Liran Einav	February 2–3, 2024
Economic Fluctuations and Growth Program Meeting	Gabriel Chodorow-Reich and Ayşegül Şahin	February 23, 2024
Mentoring Program on Aging and Health Economics Research – Virtual Information Session	Jetson Leder-Luis and Sebastian Tello-Trillo	February 26, 2024
NBER Digital Economics and AI Meeting	Avi Goldfarb, John J. Horton, Abhishek Nagaraj, and Catherine Tucker	February 29–March 1, 2024
Monetary Economics Program Meeting	Killian Huber and Eric R. Sims	March 1, 2024
TRIO Conference	Shin-ichi Fukuda, Joshua K. Hausman, and Kenichi Ueda	March 2–3, 2024
The Economics of Firearm Markets, Crime, and Gun Violence	Marcella Alsan, Philip J. Cook, and Sara B. Heller	March 7, 2024

Title of Conference/Meeting	Organizers	Dates
Environment and Energy Economics Program Meeting	Namrata Kala and Wolfram Schlenker	March 7–8, 2024
Law and Economics Program Meeting	Christine Jolls	March 8, 2024
Financial Market Frictions and Systemic Risks	Wenxin Du, Alp Simsek, and Chester S. Spatt	March 8, 2024
Immigrants and the US Economy	Aimee Chin and Kalena Cortes	March 8, 2024
Economics of Decarbonizing the Built Environment	Peter Christensen, Meredith Fowlie, and Christopher R. Knittel	March 14, 2024
CRIW Race, Ethnicity, and Economic Statistics for the 21st Century	Randall Akee, Lawrence F. Katz, and Mark Loewenstein	March 14–15, 2024
Labor Studies Program Meeting	David Autor and Alexandre Mas	March 21–22, 2024
Chinese Economy Working Group Meeting	Nancy Qian, Shang-Jin Wei, and Daniel Xu	March 29, 2024
International Finance and Macroeconomics Program Meeting	Yan Bai and Anusha Chari	March 29, 2024
Economics of Aging Program Meeting	Kathleen M. McGarry and Jonathan S. Skinner	March 29, 2024

Entrepreneurship and Innovation Policy and the Economy, Volume 3

Benjamin Jones and Josh Lerner, editors

Entrepreneurship and Innovation Policy and the Economy, Volume 3 synthesizes key findings about entrepreneurial and innovative activity in the US economy, conveying insights on contemporary challenges and providing an analytical base for policy design.

In the first paper, Jorge Guzman, Fiona Murray, Scott Stern, and Heidi Williams examine regional innovation engines and highlight the place-specific actions, potential bottlenecks, and roles of different stakeholders in catalyzing entrepreneurship and innovation.

Next, Lee Branstetter and Guangwei Li examine the challenges faced by the Chinese central government in implementing industrial policy to push the technology frontier while local governments and businesses deploy resources to advance their own, not necessarily aligned, interests.

Turning to climate issues, James Sallee analyzes policies aimed at accelerating the energy transition by hastening the replacement of durable capital assets like automobiles and residential appliances that last for decades and slow the adoption of cleaner technologies.

Joshua Gans studies cryptocurrencies and other crypto-token based instruments and the broad range of government responses to them, particularly in the US.

Finally, Ina Ganguli and Fabian Waldinger consider the effects of the Russian invasion of Ukraine on the human capital in the Ukrainian science community.

Environmental and Energy Policy and the Economy, Volume 5

Matthew J. Kotchen, Tatyana Deryugina, and Catherine D. Wolfram, editors

This volume presents six new papers on environmental and energy economics and policy. Sarah Armitage, Noël Bakhtian, and Adam Jaffe review the literature on innovation market failures with an eye towards developing insights on the implementation of such policies in the climate and energy context.

Richard Newell, William Pizer, and Brian Prest discuss alternative ways of accounting for capital displacement in benefit-cost analysis.

Tihitina Andarge, Yongjie Ji, Bonnie Keeler, David Keiser, and Conor McKenzie provide new estimates of the distribution of environmental benefits and burdens of the Clean Water Act.

E. Mark Curtis, Layla O'Kane, and Jisung Park examine the employment transitions into and out of sectors most likely affected by decarbonization.

Lucas Davis provides a detailed analysis of heat pump adoption in the United States, showing that it may be one of the few energy-efficiency technologies for which subsidy take-up does not favor high-income households.

Finally, Robert Huang and Matthew Kahn contribute to the political economy of US energy policy, showing that many Republican-leaning states have a comparative advantage at generating some types of green power.

The Economics of Artificial Intelligence: Health Care Challenges

Ajay Agrawal, Joshua Gans, Avi Goldfarb and Catherine Tucker, editors

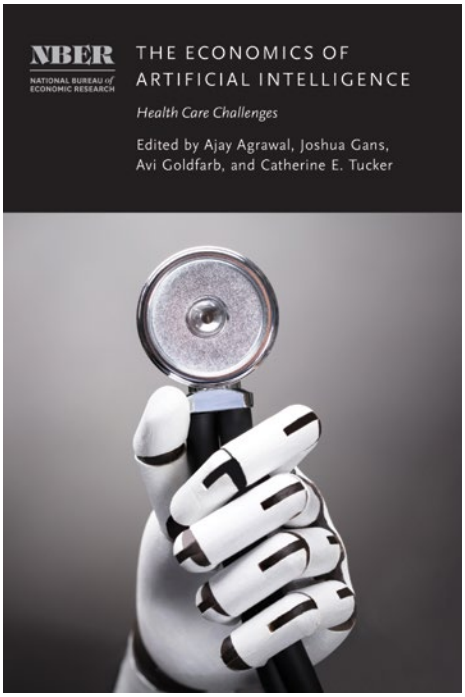

In sweeping conversations about the impact of artificial intelligence on many sectors of the economy, healthcare has received relatively little attention. Yet it seems unlikely that an industry that represents nearly one fifth of the economy could escape the efficiency- and cost-driven disruptions of AI.

The Economics of Artificial Intelligence: Health Care Challenges brings together contributions from health economists, physicians, philosophers, and scholars in law, public health, and machine learning to identify the primary barriers to entry for AI in the healthcare sector. Across original papers and in wide-ranging responses, the contributors analyze barriers of four types: incentives, management, data availability, regulation.

They also suggest that AI has the potential to improve outcomes and lower costs. Understanding both the benefits of and barriers to AI adoption is essential for designing policies that will affect the evolution of the healthcare system.

NBER | NATIONAL BUREAU of ECONOMIC RESEARCH

1050 Massachusetts Avenue
Cambridge, Massachusetts 02138-5398
(617) 868-3900

Change Service Requested