The NBER Digest

NATIONAL BUREAU OF ECONOMIC RESEARCH

September 2023

INSIDE THIS ISSUE

Unpacking the Causes of Pandemic-Era Inflation in the US

After several decades of relatively low and stable inflation, in 2021 the US experienced a sharp rise in the pace of price increases. The annual inflation rate, as measured by the Consumer Price Index, was 1.7 percent in February 2021 but rose to more than 5 percent in June 2021. It continued rising for another year, peaking at about 9 percent in June 2022.

Rising commodity prices and supply chain disruptions were the principal triggers of the recent burst of inflation. But, as these factors have faded, tight labor markets and wage pressures are becoming the main drivers of the lower, but still elevated, rate of price increase.

The rise in the inflation rate has been attributed to many factors. The US response to the COVID-19 pandemic included a series of federal initiatives, notably the CARES Act and the American Rescue Plan, which collectively authorized roughly $5 trillion in government spending. These programs contributed to strong consumer and business demand, which tightened labor markets (between mid 2021 and early 2022 the ratio of job vacancies to unemployed workers doubled), putting upward pressure on wages and prices.

On the supply side, supply chain disruptions had an important inflationary impact, particularly in 2021 and 2022. The auto industry is a case in point. US auto production dropped from 11.7 million vehicles in July 2020, roughly the pre-pandemic rate, to less than 9 million in the fall of 2021, reflecting shortages of computer chips and other inputs. The combination of strong demand and supply chain bottlenecks led to further pressure on prices, particularly on prices of durable goods. Rising prices of food and energy added importantly to inflation. Notably, the crude oil market was disrupted by the Russian invasion of Ukraine in early 2022. The price of West Texas Intermediate crude oil rose from less than $70 per barrel in the late summer of 2021 to more than $100 per barrel for most of the period between March and July of 2022, pushing up gasoline prices and the costs of many industrial inputs.

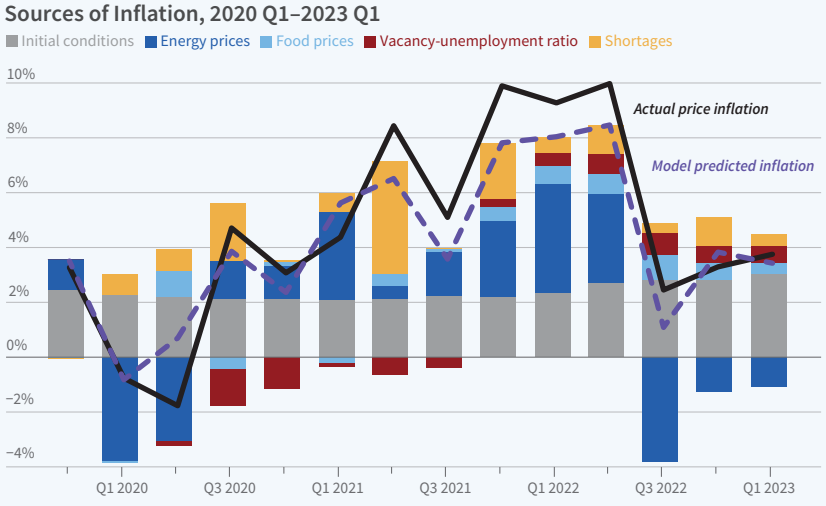

"Initial conditions" show the contribution of pre-2020 data and include the contributions of productivity shocks
Source: Researchers' calculations using data from the Bureau of Labor Statistics and the Federal Reserve Bank of St. Louis

The relative importance of these factors, and others, in contributing to US inflation that began in mid-2021 is an open question. In What Caused the US Pandemic-Era Inflation? (NBER Working Paper 31417), Olivier J. Blanchard and Ben S. Bernanke study the historical comovement of wages, prices,

and inflation expectations in an effort to measure the relative contributions of the sources of the recent US inflation shock. They estimate the relationships between price inflation, wage inflation, commodity price shocks, shortages, and labor market tightness over the period from 1990 to the start of the pandemic, and then use their estimates to simulate the inflationary effects of the various shocks that buffeted the US economy from the beginning of 2020 to early 2023.

The researchers find that energy prices, food prices, and price spikes due to shortages were the dominant drivers of inflation in its early stages, although the second-round effects of these factors, directly through their effects on other prices or indirectly through higher inflation expectations and wage bargaining, were limited. The contribution of tight labor markets to inflation w as i nitially q uite m odest. B ut as product market shocks have faded, the tight labor market and the resulting persistence in nominal wage increases have become the main factors behind wage and price inflation. This source of inflation is unlikely to recede without macroeconomic policy intervention.

—Leonardo Vasquez

Correcting for Quality Change When Measuring Inflation

One of the perennial challenges of constructing price indices like the Consumer Price Index (CPI) is that products change over time. This is often cited as a concern with regard to rapidly evolving products on the technological frontier, such as personal computers, cellphones, and automobiles. One standard approach to adjusting for quality change, the "hedonic method," involves relating product prices to product characteristics, such as memory size and CPU speed for computers or horsepower and miles per gallon for cars, and estimating the amount that consumers are prepared to pay for improvements in these characteristics. These estimates can in turn be used to distinguish price changes over time that are due to changes in product characteristics from changes that are due to the price charged for a product with a given bundle of characteristics. Price changes for the same product over time are relevant for inflation calculations; price changes that result from changes in product quality are not.

Applying the hedonic method can be difficult because most products have many different characteristics, making it difficult to choose the set of characteristics that should be included in the hedonic model, and the set of pricing-relevant characteristics may change over time. The Bureau of Labor Statistics currently uses hedonic adjustment for only about 7.5 percent of the goods that are included in the calculation of the CPI.

Analysis of changing food product characteristics suggests that conventional measures may substantially overstate food-price inflation.

In Using Machine Learning to Construct Hedonic Price Indices (NBER Working Paper 31315), Michael Cafarella, Gabriel Ehrlich, Tian Gao, John C. Haltiwanger, Matthew D. Shapiro, and Laura Zhao describe a way to substitute a machine-learning algorithm that can search for patterns and relationships in large datasets for the human analysts who traditionally construct the hedonic functions. This algorithm provides estimates of the relationship between changes in prices and product characteristics that can be used to construct hedonic-adjusted measures of inflation.

The researchers apply their inflation-measurement procedure to the Nielsen Retail Scanner data from 2006 to 2015. The data include 2.6 million products defined by 12-digit universal product codes (UPC) that uniquely identify specific goods. Each UPC has a text description that is a combination of Standard English words and industry-specific character combinations. The data are collected from over 40,000 stores from approximately 90 retail chains, representing 54 percent of all grocery store sales, 55 percent of all drugstore sales, 32 percent of all mass merchandise sales, and 2 percent of all convenience store sales. The researchers aggregate weekly data into quarterly data and normalize product quantities.

The hedonic model

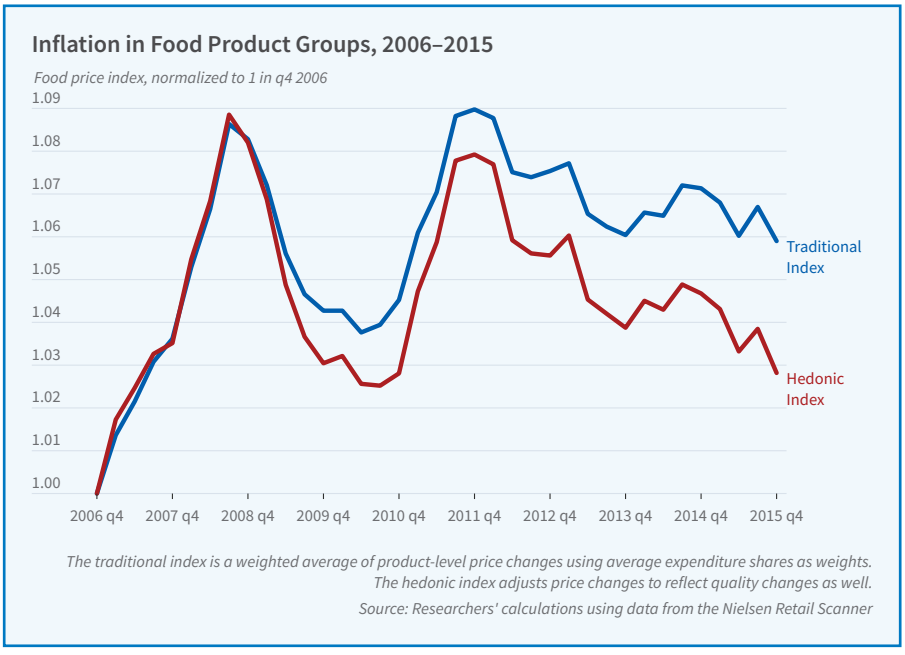

generated by the machine-learning algorithm has excellent predictive power for price changes on average, though its ability to predict price changes across products is inconsistent. Out of sample, the algorithm correctly predicts about 60 percent of the range of price changes for the median food product and about 50 percent for the median nonfood product.

The researchers use their method to adjust a traditional inflation measure to reflect a larger range of product quality changes over the 2006 to 2015 period. They find that while the standard method yields a 5.9 percent increase in food prices over the period, correcting for quality change reduces this increase to 2.8 percent. For nonfood products sold at grocery stores and drugstores, which may be less representative of products in the overall economy, the paper finds little evidence of quality change.

— Whitney Zhang

The Effects of Job Loss on Low-Wage Workers

The effect of job losses on workers' subsequent earnings is a subject of long-standing interest in labor economics. For a well-paid worker, losing a job can trigger a fall down the "job ladder" into a lower-paying position, firm, or industry. But what is the experience of low-wage workers, who are already on the bottom rungs of the ladder and have less room to fall?

In How Replaceable Is a Low-Wage Job? (NBER Working Paper 31447), Evan Rose and Yotam Shem-Tov study this question by tracking workers who earned less than $15 an hour in 2020 dollars between 2001 and 2014. They compare workers who leave their jobs because their firm experienced a large employment contraction to workers in comparable firms that did not experience a contraction. Focusing on company-wide layoffs is a standard strategy in the study of job displacement since it isolates job separations due to layoffs and reduces the chances that the workers who lose their jobs are different from those who keep their jobs in some unobserved way that matters for future earnings.

Workers who lose a low-wage job suffer from persistent earnings declines. After six years, their quarterly earnings average was 13 percent lower than that of the comparison group. Very little of this loss in earnings comes from reductions in hourly wages since these workers were at the bottom of the wage distribution prior to their job loss. Instead, about half of the long-run loss arises because displaced workers are more likely to be without a job for long stretches of time. Four to six years out, unemployment rates among displaced workers are 3.2 percentage points higher, and rates of labor force nonparticipation are 2.6 percentage points higher, than in the comparison group. The other half of the long-run earnings loss comes from a decline in the number of weeks worked per year — down by about 3.2 — and in hours worked per week — down by about 3.

> Six years after a low-wage worker loses a job, on average that worker's earnings are significantly lower than the earnings of a comparable worker without a job loss.

When the researchers apply a similar empirical strategy to an alternative sample of workers earning $15 to $30, they find long-run earnings losses of 17 percent. As in prior studies, the earnings losses for these workers come predominantly from drops in hourly wages rather than reductions in employment rates or weekly hours. Job loss therefore seems to have large effects on workers across the wage distribution, but the sources of low-wage workers' losses are different. Their difficulty seems to be climbing back onto the job ladder at all, rather than falling down to a lower rung where jobs pay less per hour.

— Shakked Noy

Earnings Trajectories Following Job Losses for Low-Wage Workers

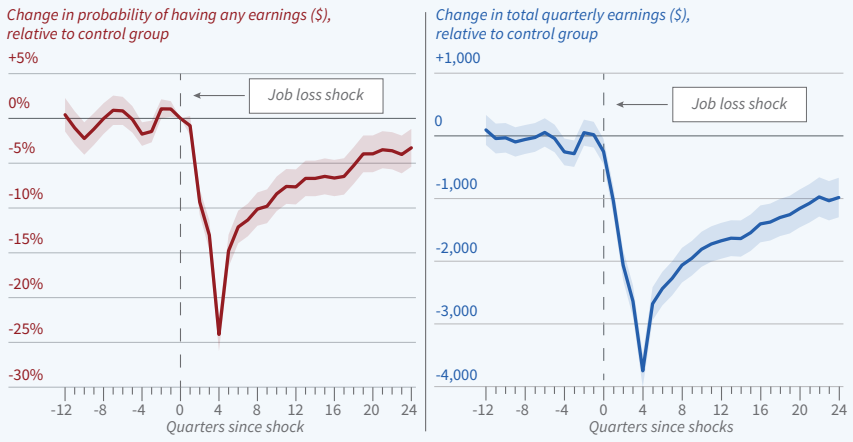

Full-time workers earning $15 per hour or less are defined as "low-wage". Job losses defined as zero earnings in a quarter from primary employer. Source: Researchers' calculations using data from the US Census

Designing a Public Transit Network: Evidence from Jakarta

TransJakarta, the bus system serving greater Jakarta, Indonesia, undertook a major expansion in the last decade. It tripled its routes and doubled the number of buses in operation between January 2016 and February 2020. In Optimal Public Transportation Networks: Evidence from the World's Largest Bus Rapid Transit System in Jakarta (NBER Working Paper 31369), Gabriel Kreindler, Arya Gaduh, Tilman Graff, Rema Hanna, and Benjamin A. Olken examine the effects of this expansion on ridership and aggregate trip flows across this 120-mile network.

The TransJakarta system integrates Bus Rapid Transit (BRT) lines that operate in dedicated bus lanes with designated stations, and non-BRT routes that run on some city streets as well as in BRT corridors, making stops at BRT stations.

The researchers study the impact of introducing new routes on bus ridership and on aggregate trip flows, measured using anonymized smartphone location data. They measure the impact of three types of changes induced by new route launches. First, they find a 16 percent increase in ridership when two locations went from being connected using transfers to being directly connected, in those cases when the new route had a similar travel time as previous connections. Second, the effect on ridership was larger — 27 percent — when the new direct route was also faster than the existing transfer connections. Third, ridership increased by 9 percent when a new direct route launch increased the frequency of buses between two locations that were already directly connected. The last effect implies that a 10 percent decrease in wait times leads to a 2.9 percent increase in ridership on BRT routes, while for non-BRT routes this elasticity is 1.05. The researchers did not find aggregate trip volume increases after new route launches.

> Bus ridership responds to improvements in wait time, travel time, and access to direct connections.

To interpret these empirical findings, the researchers develop a model of demand for public transportation and estimate it by matching the impacts on bus ridership. They find that commuters in Jakarta are 2.4 times more sensitive to wait time for the bus than time spent on the bus. This may capture the uncertainty in wait times and differences in the experience of waiting compared to riding the bus. The pure transfer penalty, which captures dislike for transfers aside from the travel and wait time involved, is small and not statistically significantly different from zero.

The researchers use the estimated demand model to compare the post-expansion TransJakarta network with networks that deliver high levels of total rider welfare. To do so, they introduce a framework for describing the characteristics of optimal networks. The current network is characterized by a concentration on the city's urban core, with relatively few bus routes connecting to the city's outskirts. In contrast, the optimal network is less dense and extends beyond the center of the city, connecting 57 percent more locations. In the optimal network, 39 percent of bus stops are connected by either a direct or a transfer bus connection; in the current network, only 12 percent are. The researchers estimate that the equivalent variation benefit of shifting to the optimal network would be shaving 23 minutes off the travel time for each bus user in the current network. Despite the high value that commuters place on wait time, these results suggest that system ridership would have been higher if the network had been focused on broad expansion rather than on intensified service in the urban core.

—Leonardo Vasquez

Current TransJakarta Bus Network Versus an Optimal Network

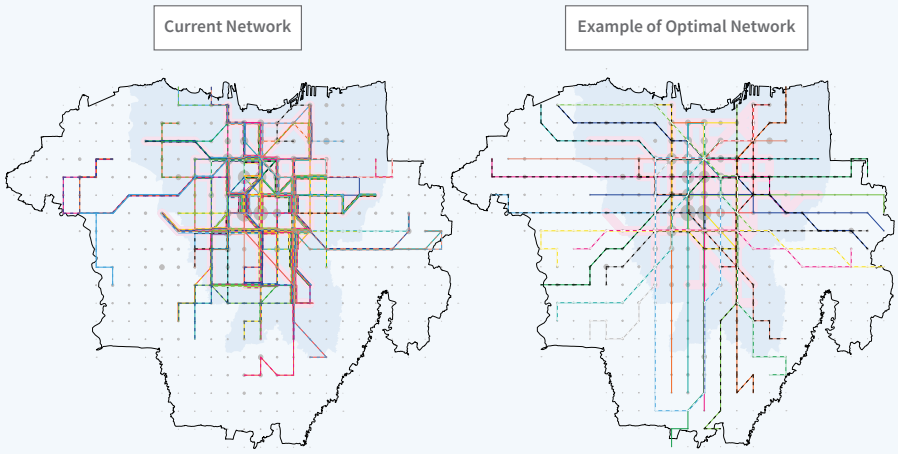

Shaded blue area represents the central urban Jakarta area
Source: Researchers' calculations using administrative ridership, anonymized smartphone location, and bus location data

How Does ESG Investing Affect an Institution's Portfolio Composition?

Financial institutions that mention environmental, social, and governance (ESG) criteria in their investment policies had $35 trillion in assets under management in 2020. To assess how these criteria affect portfolio composition, however, it is necessary to compare their actual portfolio holdings with a counterfactual that describes what their portfolio holdings would have been in the absence of any ESG considerations. In Green Tilts (NBER Working Paper 31320), Lubos Pastor, Robert Stambaugh, and Lucian Taylor estimate the extent to which ESG factors alter investment portfolios.

The researchers first note problems with the usual measure of the amount of ESG investment, which is simply the total assets of institutions declaring an ESG objective. Those institutions hold varying mixes of "green" (ESG-favored) and "brown" (disfavored) stocks. Counting *all* of their assets as green overestimates the importance of ESG investing. At the same time, institutions that do not explicitly mention ESG criteria in their investment policies might engage in ESG investing. For example, they might expect higher returns on ESG stocks than on other stocks. This would create a downward bias in the estimated importance of ESG considerations.

To estimate the importance of ESG investing, the researchers analyze institutional investors' holdings of US stocks as reported in SEC form 13F filings. The sample includes US-based institutions with more than $100 million in stock investments — a group that includes investment companies, banks, insurance companies, pension funds, and endowments — as well as foreign institutions that hold US stocks.

For each stock held by each institution, the researchers estimate a portfolio allocation share absent any ESG considerations. They define the difference between the observed share and the predicted share absent ESG considerations as the stock's "ESG tilt" for this institution. They further decompose this tilt into an extensive-margin component, reflecting the institution's binary decision about whether or not to hold a particular stock, and an intensive-margin tilt, reflecting the amount of the stock held conditional on a positive holding. Intensive-margin tilts are twice as important as extensive ones in their analysis. The researchers aggregate the ESG tilts for specific stocks and institutions to create institution- and market-level estimates of the extent to which ESG considerations affect portfolio composition.

The results suggest that about 6 percent of US equity investment allocations represent ESG tilt, a much lower figure than the percentage of total assets that are held by institutions that subscribe to ESG principles. The aggregate importance of ESG tilts is relatively constant over the 2012–21 period, reflecting a combination of increasing ESG tilts in institutions' actively managed portfolios, from about 14 percent to 25 percent over the last five years, and a decline in active management by institutional investors over the same period. Passive strategies, such as investing in all of the stocks in a stock index, do not involve any ESG tilt.

> About 6 percent of US equity investments are explained by "ESG tilt," much less than the total assets managed by institutions subscribing to ESG principles.

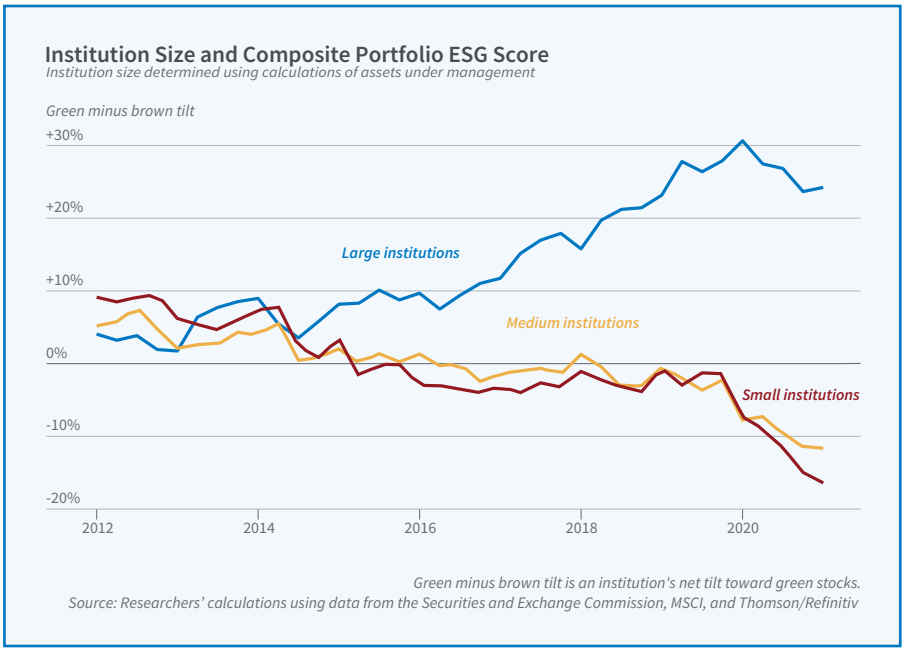

The researchers' ESG tilt measure reflects the share of an institution's investment choices that can be explained by stocks' ESG characteristics, which captures choices that both actively favor green stocks and disfavor brown stocks. The two can be separated. Within the largest firms in the investment industry, green tilts dominate brown tilts, and the strength of green tilts is growing over time. Most of the reductions in brown-stock holdings occur on the intensive margin, rather than as complete divestment.

The extent of green tilt is positively related to the assets under management of the 13F-filing institution. It is also greater for institutions — accounting for 76 percent of the assets under management in the sample — that have signed the United Nations' Principles for Responsible Investment.

—Shakked Noy

Contractor Incentives Improve Performance of Weatherization Program

In a program designed to save energy by promoting home weatherization among low-income households, compensating contractors based on the amount of natural gas saved increased the average benefit by 24 percent. Peter Christensen, Paul Francisco, and Erica Myers report this finding in Incentive Pay and Social Returns to Worker Effort in Public Programs: Evidence from the Weatherization Assistance Program (NBER Working Paper 31322).

The researchers partnered with the Illinois Home Weatherization Assistance Program and studied projects that were undertaken in 2018 and 2019. Jobs were randomly assigned contracts of three types: high incentive, low incentive, and as a control, no incentive. Projects in Cook County, which includes Chicago, were excluded because the city had a performance-based pay system already in place.

Under the program, an energy specialist establishes a baseline for each project by installing a large fan in the frame of one of the home's outside doorways. The fan draws air out of the home, reducing the air pressure inside. Leakier homes require that more air be sucked out to achieve the same amount of depressurization. Based on these test results, the weatherization program calculates the target for the potential benefits of air sealing.

> Offering bonuses to contractors who exceeded a target effectiveness level resulted in improved performance whether the bonuses were large or small.

The extent to which a home is "sealed" against the weather depends on its condition and the skill and motivation of contractors to seek out leaks in attics, walls, basements, chimneys, crawl spaces, and other places. Bonuses for each project were tied to the amount by which contractors exceeded the minimum leakage reduction goal for the home.

Although jobs awarded the higher bonuses achieved better results than those awarded smaller ones, the difference between the two groups was not statistically significant. Both large and small bonuses decreased the probability of a callback, when the contractor must correct deficiencies, by about a third. The higher-bonus group appeared to have fewer deficiencies than the lower-bonus one, but once again the difference was not statistically significant.

Using gas bills, the researchers estimate that pay-for-performance incentives for contractors reduced gas consumption by 20 to 25 percent. Contractors considered high quality based on pre-study performance achieved twice the reduction in leakage of their lower-quality counterparts. The researchers conclude that their low incentive option, which offers only a modest payment to contractors, yields greater returns per dollar than their higher incentive option.

— Steve Maas

Incentives and Contractor Performance in Weatherization Assistance Program

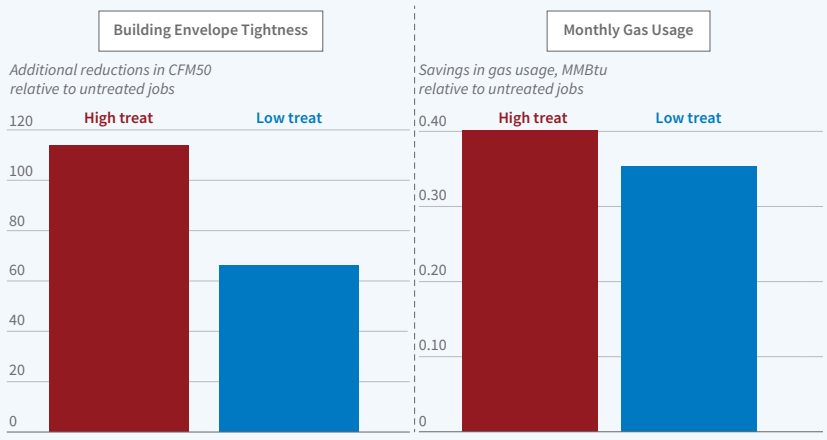

Low treat ($0.40) and high treat ($1.00) refer to piece-rate bonuses per 50 cubic feet per minute reductions.
Source: Researchers' calculations using data from the IHWAP and inspection reports from multiple agencies

NBER

The National Bureau of Economic Research is a private nonprofit research organization founded in 1920 and devoted to conducting and disseminating nonpartisan economic research. Its officers are:

James M. Poterba—President and Chief Executive Officer
John Lipsky—Chair
Peter Blair Henry—Vice Chair
Robert Mednick —Treasurer

The NBER Digest summarizes selected Working Papers recently produced as part of the NBER's program of research. Working Papers are intended to make preliminary research results available to encourage discussion and suggestions for revision. Neither the Working Papers nor The Digest have been subject to peer review or review by the NBER Board of Directors.

The Digest is free and may be reproduced with appropriate attribution to the NBER. Please provide the NBER's Public Information Department (ktasley@nber.org) with copies of anything reproduced.

Requests for Digest subscriptions, changes of address, and cancellations may be sent to Digest, NBER, 1050 Massachusetts Avenue, Cambridge, MA 02138-5398 (please include the current mailing label), or emailed to subs@nber.org. Print copies of the Digest are only mailed to subscribers in the US and Canada; those in other nations may request electronic subscriptions at www.nber.org/drsubscribe/.

Individual copies of NBER Working Papers are available online free of charge to affiliates of subscribing organizations, such as universities and colleges, and to employees of NBER corporate associates. All visitors to the NBER website receive three free downloads each year, after which there is a charge of $5 per downloaded paper. To place an order, please email the NBER's Subscriptions Department at subs@nber.org. A full subscription to the NBER Working Paper series entitles the subscriber to all new papers, recently more than 1,200 per year. The standard annual rate for a full digital subscription is $2,995; the online academic rate is $1,375. Partial subscriptions are also available; rates may be found at nber.org/wpsubscribe.html