

INSIDE THIS ISSUE

- Intergenerational Impacts of the Perry Preschool Project
- Buyer Inertia and the Cost of Health Insurance
- The Disproportionate Death of Ukrainians in the Soviet Great Famine
- Voluntary Self-Monitoring in the Auto Insurance Market
- The Race to Exploit Stock Price Differences between Exchanges

Among Umpires, Calls Improve When the Stakes Are High

Do individuals prioritize their limited supply of attention and focus more on important than on less-important decisions? In **The Dynamics of Inattention in the (Baseball) Field** (NBER Working Paper 28922) James E. Archsmith, Anthony Heyes, Matthew J. Neidell, and Bhaven N. Sampat investigate this question by analyzing the intertemporal, high-stakes decision-making of home plate umpires in Major League Baseball (MLB).

The researchers utilize data on approximately 3.2 million umpire calls during regular innings in 26,523 games over the 2008–18 period. Their study considers camera information that provides a record of the objectively correct calls in each game. They find that umpires not only allocate more attention to high-stakes calls, as indicated by the accuracy of their calls, but also that umpires devote less attention to calls that occur after they have expended a high level of attention on earlier decisions, or before they anticipate having to devote elevated attention to consequential future pitches.

MLB umpires in the researchers' sample call an average of 120 pitches each game. To

measure how pivotal each of these pitches is in influencing the outcome of a game, the researchers estimate the “leverage” of each of the calls: the absolute difference in the probability that the home team wins should the

influence on the outcome of the game.

While leverage can evolve considerably from pitch to pitch, the average impact of the call associated with any single pitch is low, implying that its potential to affect which

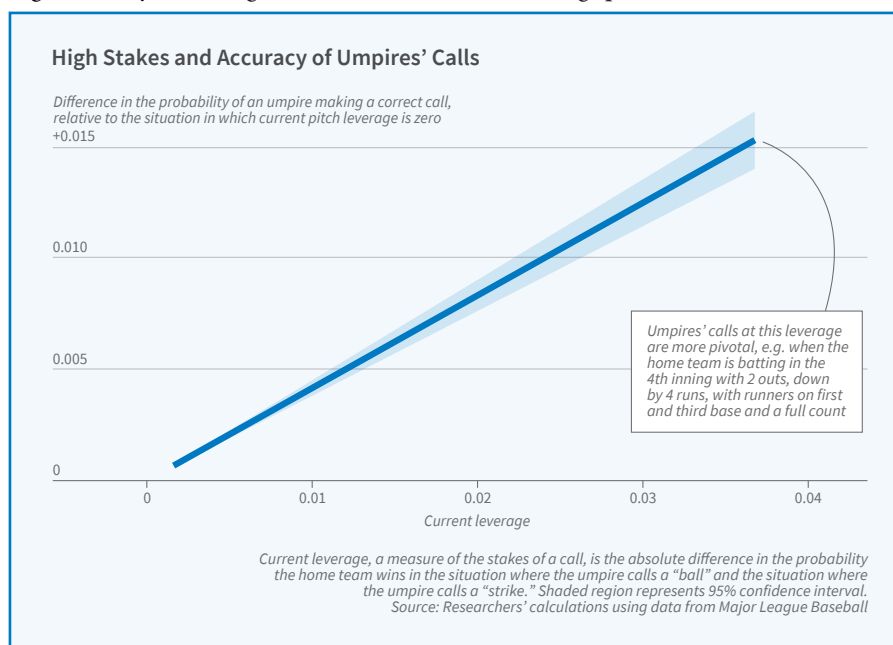
The probability of an accurate ball or strike call rises when the pitch is more likely to affect the outcome of the game. Umpires devote less attention, however, to calls made after or in anticipation of high-stakes decisions.

umpire call a strike versus a ball. For example, the call when there are no outs, no balls, two strikes, the bases loaded, and the batting team down by three runs in the top of the 9th inning, would fall in the 99th percentile of the leverage distribution because it has a large potential

team wins a game is minimal. The researchers find that an increase in the current leverage of a pitch is associated with an increase in the accuracy of an umpire's call due to higher attention paid to this high-stakes decision. A one standard deviation increase in the stakes of

the call increases the probability of an umpire making the correct call by 0.61 percent.

Increases in past leverage—the sum of the leverage of all previous calls in the present inning—also decrease the accuracy of an umpire's contemporaneous calls. A one standard deviation increase in past leverage reduces the probability of an umpire calling a pitch correctly by 0.32 percent. This is consistent with umpires having an



exhaustible supply of attention, with the attention applied to a particular decision decreasing the availability of attention for succeeding pitches. Breaks between every half-inning appear to allow umpires to regain their attention reserves, since higher leverage in preceding innings has no sizeable effect on

umpires' attention to current calls.

Allocation of attention is also shaped by expectations about the future. Expected future leverage—the sum of expected leverage of future pitches in the current inning—decreases the accuracy of an umpire's contemporaneous calls. An increase of one standard deviation

in expected future leverage decreases the probability of an umpire's correct call by 0.49 percent, consistent with umpires maintaining their limited attention stock by paying less attention to current calls when they anticipate higher-stakes pitches later in the inning.

—Aaron Metheny

Intergenerational Impacts of the Perry Preschool Project

Using newly collected late midlife measures of skills and life cycle panel data from surveys and administrative criminal records, [Jorge Luis García](#), [James J. Heckman](#), and [Victor Ronda](#) explore intergenerational impacts of the Perry Preschool Project (PPP), an early childhood education program designed to promote social mobility among disadvantaged African American children. Their study, **The Lasting Effects of Early Childhood Education on Promoting the Skills and Social Mobility of Disadvantaged African Americans** (NBER Working Paper 29057), finds that the program had multiple positive impacts, several of which affect the children of PPP participants.

Launched in Ypsilanti, Michigan in 1962, the PPP was intended to foster development of cognitive and socioemotional skills and promote social mobility. Eligibility criteria based on IQ scores and socioeconomic status were used to create a pool of 123 disadvantaged African American children who were randomized into treatment and control groups. Treatment group children received two years of two-and-a-half-hour preschool sessions on weekdays starting at age three, as well as weekly teacher home visits during the treatment period. Control group children received no interventions. The researchers analyzed the effects of the PPP because it con-

tinues to influence the design of current and proposed early childhood education programs. At least 30 percent of current Head Start programs are based on it, and more than 10 percent of African American children born in the 2010s would be eligible for it.

The program had long-lasting effects on cognition, and first-generation treatment group members also have more stable home lives and higher incomes in their child-rearing years.

Using new data on the original PPP participants, and on their children, many of whom are now in their mid-20s, the researchers find long-lasting beneficial program effects. They find positive effects on cognition for participants through age 54, contradicting claims of fadeout that have dominated popular discussions of early childhood programs. First-generation treatment group members also have more stable home lives in terms of marriage


and divorce, as well as higher incomes in their child-rearing years.

These improved outcomes across the life cycle have translated into better family environments for the second generation, the children of the PPP participants, and these improved

home environments are the source of the PPP's intergenerational benefits. Children of the first generation of treatment participants were more likely than children of the first generation of control participants to grow up in stable two-parent households. Their parents have higher average earnings, less engagement with the criminal justice system, and better executive functioning (cognition), socioemotional skills, and health. The children of treatment

group members are also less likely to be enrolled in special education, and they are 17 percentage points less likely to have been suspended from school during K-12 education than the children of control group members. They are more likely to be high school graduates and much less likely to engage in crime. They are also 11 percentage points more likely to be in good health through young adulthood, 26 percentage points more likely to be employed, and

Earnings and Marriage Outcomes of Perry Preschool Project Participants



The treatment group consists of participants in the Perry Preschool Project. Source: Researchers' calculations using data from the HighScope Educational Research Foundation

8 percentage points less likely to be divorced.

The researchers find important differences in the PPP's impact by gender. The male children of the male treatment group members

receive the greatest benefits, consistent with literature on the adverse effects of disadvantaged environments on boys. For example, children of male treated participants are 18 percentage

points less likely to have been arrested through young adulthood compared to children of male control participants.

—Lauri Scherer

Buyer Inertia and the Cost of Health Insurance

Many US markets for health insurance rely on managed competition: consumers choose freely among private insurers while regulations mandate minimum benefits and limit insurers' incentives to screen consumers by risk. Anything that impedes consumers' ability to choose high-quality plans, so-called "frictions," may undermine the efficiency of these markets.

In **Inertia, Market Power, and Adverse Selection in Health Insurance: Evidence from the ACA Exchanges** (NBER Working Paper 29097), [Evan Saltzman](#), [Ashley Swanson](#), and [Daniel Polsky](#) examine the equilibrium impact of inertia—consumers' tendency to stick with the same plan from year to year even when premiums change and better options arise—on the efficiency of California's Affordable Care Act (ACA) exchange, given the features that manage competition in this market.

The researchers analyze consumer-level enrollment data from California's ACA exchange from 2014 through 2018. Consumers choose between ACA insurance plans standardized into four tiers of increasing generosity, labeled bronze, silver, gold, and platinum. There is substantial churn from year to year in the set of consumers who are eligible to purchase insurance on the exchange. Consumers

who remain eligible and stay insured in two consecutive years make notably persistent choices: about 79 percent choose the same plan, 91 percent choose a plan from the same tier, and 87 percent choose a plan from the same insurer as in the previous year. The degree of inertia can be estimated by comparing the active enrollment

choices among those who are newly insured on the exchange, and the choices of those who were insured on the exchange in the previous year. The high degree of inertia is costly to consumers, who on average forego up to \$2,324 in annual surplus by remaining in their previous

If consumers in California's ACA exchange fully optimized their policy choices each year, average monthly premiums would be \$58 lower.

plan rather than making a new, and optimizing, choice.


To model the ACA market, the researchers assume that oligopolistic insurers set premiums to maximize their expected profit, while consumers, presented with a menu of policies, choose plans to maximize their utility. When insurers know that their buyers are slow to

cent. If they made optimal choices each year, consumers would experience large welfare gains. Some enrollees would exit the exchange, while others would shift to less costly or higher quality plans within or across the tiers. The average per capita welfare gain would be \$902.

That combines the gain to exchange participants from selecting higher-value plans, and the savings to taxpayers from reduced ACA premium subsidies.

Two key features of the ACA exchange, risk adjustment and imperfect competition, have an important effect on the potential gains to eliminating inertia. Risk adjustment is a

requirement that firms with lower-risk consumer pools must transfer money to those with higher-risk pools. Without risk adjustment, there would be a stronger positive relationship between plan generosity and premiums. The market for high-coverage insurance could unravel, and eliminating inertia could exacerbate this risk as consumers become more attentive to their premium and risk status. The researchers find



react to premium changes or other changes in market conditions, they have greater latitude to charge high prices.

Under their assumptions about market structure and the way insurers respond to inertia, the researchers estimate that eliminating inertia would decrease premiums by 13.2 per-

cent. only limited support for this possibility. Even without risk adjustment, they estimate that removing inertia would generate a per capita welfare gain of \$658, with a premium decline of 10.9 percent.

Absent market power, eliminating inertia would also have smaller welfare benefits. In the

baseline analysis, insurers exploit consumers' inertia to set marked-up premiums; eliminating inertia steeply reduces prices. When com-

petition forces insurers to price plans at average cost, however, eliminating inertia reduces premiums by only 0.5 percent. Nevertheless,

it improves consumer choice enough to raise annual average per capita welfare by \$547.

—Lucy E. Page

The Disproportionate Death of Ukrainians in the Soviet Great Famine

When harvests failed in the Soviet Union in the early 1930s, ethnic Ukrainians experienced significantly higher mortality than the majority Russian population. [Andrei Markevich](#), [Natalya Naumenko](#), and [Nancy Qian](#) discover in *The Political-Economic Causes of the Soviet Great Famine, 1932–33* (NBER Working Paper 29089) that this was due in large part to intentionally biased economic policies of the central government.

Drawing from archival sources, including the Soviet census, planning documents, and declassified secret police records, the researchers construct a dataset of demographic, economic, political, historical, geographical, and climatic factors for the years 1922–40. Their sample includes data from 19 provinces in the three most populous republics of the Soviet Union — Belarus, Russia, and Ukraine — comprising 84 percent of the Soviet population and 88 percent of the population in grain-producing areas. The researchers use this information, along with rich data on the ethnic makeup of many districts, to examine the effects of the famine and centralized planning decisions on ethnic Ukrainians and on those of other ethnicities in Ukraine and elsewhere in the Soviet Union.

In 1926, Russians comprised 53 percent of the Soviet population, while Ukrainians, the largest ethnic minority, made up 21 percent. Yet between 30 and 45 percent of the 10.8 million victims of the Soviet Great Famine were ethnic Ukrainians.

Government policies motivated by ethnic bias can account for up to 92 percent of the deaths of ethnic Ukrainians living in Ukraine at the time.

The researchers estimate that a 10 percent increase in the ethnic Ukrainian share of the population in a province was associated with a 0.51 percentage point increase in the famine mortality rate.

They find that provinces with a greater Ukrainian population share experienced a greater increase in mortality between the

side Ukraine proper.

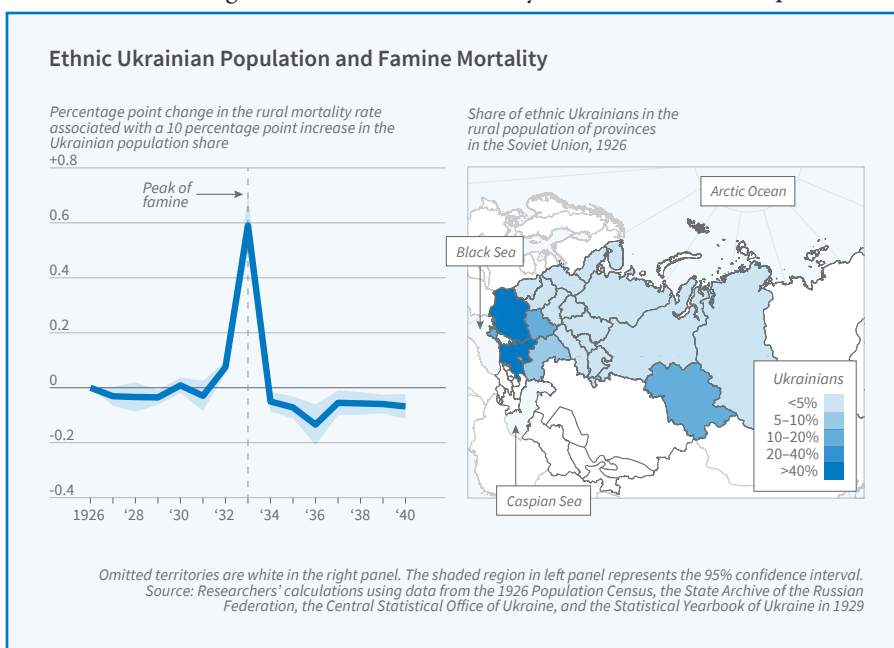
The researchers suggest that the disproportionate loss of life among Ukrainians was due to economic policies implemented by Russian Bolshevik elites and directed against ethnic Ukrainian peasants. They estimate that these policies accounted for up to 92 percent of deaths of ethnic Ukrainians

in Ukraine and up to 77 percent of deaths of ethnic Ukrainians elsewhere in the Soviet Union.

The Bolsheviks' need to control grain production led to their repressing Ukrainians more than other groups. Ethnic Ukrainian peasants played a central role in grain production. They were the largest ethnic group in designated “grain-surplus” areas, where production far exceeded subsistence levels during non-famine years.

They also had a strong group identity, a history of confrontation with the Bolsheviks during the civil war and they had resisted Soviet efforts to control agriculture, which constituted nearly half of GDP. By repressing them, Moscow advanced its efforts to control rural grain production. Centrally planned policies known to have contributed to famine mortality, such as the withholding of tractors, were more intensely enforced in regions with a greater number of ethnic Ukrainians.

—Brett M. Rhyme



pre-famine and famine years. This pattern emerges even when controlling for factors such as weather, food production, and urbanization, as well as *dekulakization* — the policy of eliminating wealthy peasants who resisted collectivization of farms — and the drop in livestock a few years prior to the famine, which could have affected grain production and the ability to survive harvest shortfalls. Further, the increase in Ukrainian mortality occurred only during the famine, and only in ethnic Ukrainian-majority areas, even those out-

Voluntary Self-Monitoring in the Auto Insurance Market

New technologies have allowed consumers to monitor their own behaviors and to sell their data to firms. In the auto insurance industry, for example, some insurers now reward customers for having devices in their cars that track their driving behavior. Insurers can then use this proprietary data to infer drivers' accident risk and to adjust their premiums accordingly.

In Buying Data from Consumers: The Impact of Monitoring Programs in US Auto Insurance (NBER Working Paper 29096), Yizhou Jin and Shoshana Vasserman examine the first major data-sharing program of this form in the US auto insurance industry. They collect detailed plan and customer data on more than one million customers during the 2012–16 period from the large insurance firm that sponsored this program.

The researchers find that on average monitoring promotes safer driving and attracts safer drivers. Consumers who opt into the program make 30 percent fewer insurance claims during the monitoring period than in subsequent periods. Moreover, drivers who select into monitoring are less risky post-monitoring, and monitoring scores strongly predict later accident claims. Monitoring thus allows insurers to collect what would otherwise be hidden information on drivers' risk types.

The researchers also investigate customers' demand for insurance and their decision to opt into monitoring based on their level of accident risk, risk aversion, expectations of future premiums, and aversion to being monitored. They estimate that on average, customers would implicitly pay \$93 to avoid monitoring, perhaps reflecting privacy concerns or higher driving effort when being monitored.

Monitoring allows the firm to offer lower

rates to safer drivers. However, since the firm sponsoring the program is the only insurance company that knows its drivers' monitoring scores, it can charge safe drivers a premium that is below the offerings of competing firms,


An opt-in program allowing drivers to provide their insurer with information on their driving behavior reduced the volunteers' premiums and raised the company's profits.

while still charging a higher markup over its true costs. In contrast, competing insurers lose their safer customers to the monitoring firm, and their profits decline. Overall, the researchers conclude that the sum of welfare gains to consumers, increased profits for the monitoring firm, and lower profits for rival insurers is positive—about \$13 per driver per year.

later be recouped. To capture this dynamic, the researchers model a pricing strategy in two phases. In the first, the firm invests in attracting consumers to the monitoring pool, either by offering an upfront dis-

count to those who opt in or by surcharging those who opt out. The researchers calculate that the optimal discount would be higher and would attract many more customers to monitoring than the status quo.

In the second stage, once information on driver behavior has been “harvested,” the insurer uses the data to set more profitable renewal premiums. The researchers conclude that if the firm made optimal strategic use of the information that it collected, it would share less of the potential premium reduction with consumers than it did in practice. Profit-maximizing pricing—both a steeper discount for being monitored, which would lead to more drivers being monitored, and a smaller premium reduction for demonstrated safety—would increase the overall gains from the monitoring program



About two-thirds of this gain is due to monitoring inducing safer driving; the remainder is due to the fact that safer drivers can purchase insurance that is priced more accurately for their risk type.

The researchers also ask whether the firm could profit by adjusting the pricing strategy that is observed in the data. A profit-maximizing strategy would balance two opposing incentives: drivers need substantial discounts at sign-up in order to opt into monitoring, but these discounts must

from about \$13 to about \$21 per driver.

If proprietary ownership of monitoring data were banned, the incentives for insurance firms to invest in monitoring programs would be severely reduced because competitors could free-ride on the data that is collected. The researchers conclude that while this would lead to lower premiums for monitored consumers, the welfare gains would be more than offset by losses from firms choosing to invest less in monitoring in the first place.

—Lucy E. Page

The Race to Exploit Stock Price Differences between Exchanges

Reducing latency, the amount of time it takes an electronic order to reach a trading venue, has become increasingly important to traders looking to profit from small, temporary price differences between exchanges. Latency-arbitrage competitions are measured in microseconds (millionths of a second), and high-frequency trading firms have invested in early access to price data and in ways to expedite orders, such as locating their computers next to the exchange's computers in the exchange's data center. Systematic empirical analysis of the scope and market implications of latency arbitrage is difficult, however, because conventional limit-order-book data do not capture failed trades and cancels, which is like only seeing the winner of a race and not the losers—the swift, but comparatively slower traders. This makes it hard to know for sure that a race happened.

Matteo Aquilina, Eric Budish, and Peter O'Neill present new evidence on this issue in **Quantifying the High-Frequency Trading “Arms Race”** (NBER Working Paper 29011). They utilize message data from the London Stock Exchange for all stocks in the FTSE 350 Index for 43 trading days in the fall of 2015. Their sample of 2.2 billion messages sent between firms and the exchange includes the failed limit or cancel order attempts of market participants, which makes it possible to directly examine races

between high-frequency traders. The researchers find that latency-arbitrage races are extremely fast and that, while individual races are not particularly lucrative for the traders or consequential for the market, the large number of races leads to

Latency-arbitrage races constitute about a fifth of total daily FTSE 350 Index trading volume and cumulate to about \$5 billion annually in global equity markets.

a significant loss of market liquidity.

The researchers define the length of a latency-arbitrage race as the time difference between the first message received by the exchange that results in a successful trade or cancel and the first message that is a failed trade or cancel. They define those whose orders are successfully executed as the winners. These competitions happen incredibly quickly. The average race takes 79 microseconds (0.000079 seconds), and modal races last between 5 and

10 microseconds. Races occur frequently; on average, there are 71,000 daily races in the FTSE 350 and one race per minute per FTSE 100 ticker. Latency-arbitrage races constitute about 20 percent of total daily trading volume.

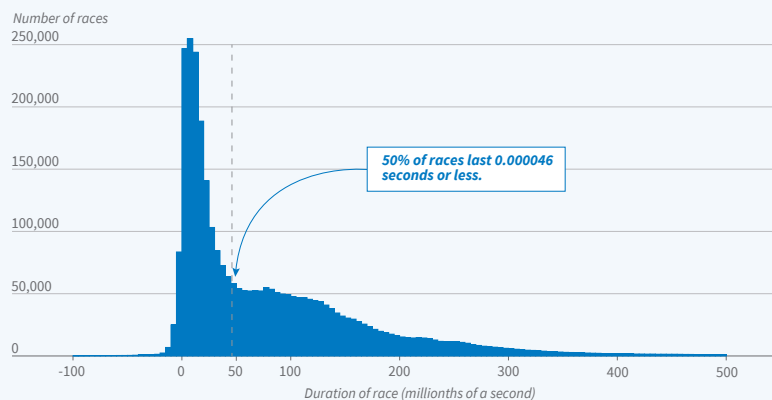
The profit from an average competition is small, about two pounds sterling, but because there are many races, the aggregate profits are substantial. The researchers calculate the average “tax” on liquidity associated with latency-arbitrage races, the ratio of daily race profits to trading volume, to be approximately 0.53 basis points. This hit to liquidity is mostly driven by the top six high-frequency trading firms “sniping” stale quotes: quickly purchasing assets with outdated prices ahead of other traders to immediately sell them for profit.

Such actions reduce, rather than increase, liquidity in the market. The top firms on net remove more liquidity from the market than they provide.

The researchers estimate that latency-arbitrage races account for about one-third of the price impact of assets involved in these trades, and that financial market reforms designed to remove latency arbitrage could decrease investor liquidity costs by 17 percent.

—Aaron Metheny

Speed of Latency-Arbitrage Races in High-Frequency Trading



The duration of a race is the time difference between the first message received by the exchange that leads to a successful trade or cancel and the first message that leads to a failed trade or cancel. This difference can be negative because of a small amount of randomness in the exchange system architecture. Source: Researchers' calculations using data from the London Stock Exchange

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. It is not copyrighted and may be reproduced with appropriate attribution of source. Please provide the NBER's Public Information Department (caradin@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/drsunsubscribe/.

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 or call (617) 588-1405. 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 \$2,675; the online academic rate is \$1,230. Hard-copy subscriptions and partial subscriptions also are available; rates may be found at nber.org/wpsubscribe.html.