

Editors' Summary

THE BROOKINGS PANEL ON Economic Activity held its eightieth conference in Washington, D.C., on September 8 and 9, 2005. This issue of the Brookings Papers includes the papers and discussions presented at the conference. The first paper examines some possible drawbacks to greater transparency by central banks. The second analyzes the change in the income distribution that has accompanied the productivity surge of the past decade. The third studies how workers respond to information about risk in their retirement accounts. The fourth looks at the long-run interactions between energy needs and the environment. The issue concludes with a report that questions the widespread preference for the payroll data over the household data as a measure of monthly employment gains.

CENTRAL BANKERS AND ECONOMISTS alike have increasingly come to believe that central bank communication and transparency are key to an effective monetary policy. For many this belief stems from the generally held view that long-term interest rates, which the central bank does not directly control, are the main link between monetary policy and the spending decisions of households and firms. Expectations about the path of future short-term rates, which the central bank does control, play a central role in determining long-term rates, and transparency regarding the setting of short-term rates is seen as increasing the central bank's ability to affect those expectations. These beliefs have led to increased openness in the conduct of policy. In 1994 the Federal Reserve's main policymaking body, the Federal Open Market Committee (FOMC), began the practice of immediately following each of its meetings with a public announcement, and the information content of those announcements has increased steadily since then. In the current recovery, for example, the FOMC's postmeeting announcements have indicated clearly the committee's intention to raise the federal funds rate, its principal short-term policy rate, at a measured pace

over a sustained period. Accompanying these announcements has been a large increase in the number of public speeches by FOMC members, aimed at better informing the public on how the committee makes its judgments. Meanwhile the central banks of many countries other than the United States have adopted explicit inflation targeting. This policy is motivated by the same belief that the central bank should be more transparent and accountable to the public. In the first article of this issue, Stephen Morris and Hyun Song Shin analyze the subtle ways in which increased transparency affects public expectations and the information value of market signals. Although they support the notion that clear statements both of the central bank's view of the economy and of its policy intentions will move expectations closer to the bank's own view, they argue that, paradoxically, such statements may actually decrease the information content of market signals, to the detriment of both the central bank and private agents.

The authors begin with an extensive discussion of the importance of informational efficiency to economic welfare. The efficacy of decisions by both private agents and the central bank depends on the accuracy of the information they have available. For example, informational efficiency plays an important role in determining the accuracy with which prices in financial markets reflect both business opportunities and resource scarcities. For central banks the primary issue is how transparency affects the prices relevant to the allocation of capital across time, rather than across firms or industries at a given moment in time. The authors describe the importance of the yield curve in affecting the duration of investment projects, the effect of long-term rates on the housing market, and the sensitivity of investment to stock prices.

Efficient information is important not just for the signals it provides to private agents, but also for the guidance it provides for the central bank's own decisions. The authors observe that, in order for the central bank to steer the economy in the right direction, it must have good information on the current state of the economy—in particular, how close the economy is to capacity—and on its likely course in the future. This includes not just information about asset prices in financial markets, but also information about goods prices and the cost conditions of firms.

An essential element in the authors' argument, not usually mentioned in discussions of central bank transparency, is that private agents collectively have a great deal of information that is relevant to efficient allocation and that, with more public information available, less of that private infor-

mation may be revealed. The authors revisit Friedrich von Hayek's classic 1945 essay, "On the Use of Knowledge in Society," which emphasizes the dispersion of information among individuals and the role of market prices in aggregating that information. For example, an individual business executive is likely to have the best estimate of the expected productivity of his or her own various investment opportunities, but not those of other businesses. If a central planner knew the complete list of potential projects of all businesses and their expected productivities, it would be a relatively simple task to allocate aggregate saving among those projects. But planners do not have that information. In Hayek's view, prices emerging from markets do convey that information, efficiently summarizing the shifting fundamentals of the economy that no individual agent can discern. Morris and Shin see the position of the central bank as analogous to that of Hayek's central planner: "The central bank's resources and expertise, as formidable as they are, may fail to match the collective wisdom of the economically active population as a whole."

If the public information provided by the central bank simply added to the private information of agents, it would be hard to imagine it being detrimental to efficiency—the more information, the better. But more public information may instead substitute for private information, by leading market participants to ignore or downplay their own private information in making their decisions. Increased central bank information might even reduce market participants' efforts to obtain information on their own. The authors describe this as an informational externality generated by public signals. Further, they stress that this externality can be exacerbated if agents have an interest in the beliefs and resulting actions of others.

To illustrate this phenomenon, they cite Keynes's analogy of financial markets to the newspaper "beauty contests" of his day, in which readers won not by picking the face that they themselves thought most beautiful, but by guessing which face would be chosen as the most beautiful by the most other readers. As in the beauty contest, Keynes argued, an investor buying stocks can win by anticipating which stocks will become the popular choice, which are not necessarily those of the companies with the best fundamentals. The authors observe that in such situations the pronouncements of the central bank about the economic outlook "provide a powerful rallying point around which market expectations can coalesce," and that "The more market participants are concerned with the beliefs of other market participants, the greater will be the impact of the central bank's pro-

nouncements in determining the aggregate market outcome.” The authors believe that, however informative the central bank’s statements may be about the fundamentals, they are even better as guides to what other agents will think. Hence traders will give them undue weight and place less weight on their own assessments of the fundamentals. The authors conclude that “Public pronouncements can thus crowd out private opinions, and the market may cease to function as a way of aggregating and revealing diverse, private judgments about the world in the way that Hayek envisaged.” This argument applies as well to increases in the accuracy of the central bank’s forecasts. Improving their accuracy increases both the weight that a typical agent will place on them and his or her belief about the weight that other agents will assign to them. It is even possible that improvement in central bank transparency or in the accuracy of its forecasts will crowd out more information than it adds, decreasing overall informational efficiency.

It is difficult to verify empirically the importance of these considerations: transparency is hard to quantify, either in levels or in changes, and it is harder still to measure “informational efficiency.” But the authors report some evidence that supports certain elements of the story. Most observers of financial markets would agree that central bank statements about the economy can have significant effects on those markets, at least in the short run, and that transparency about which economic variables the central bank is watching increases the salience of those variables in markets. The authors give two examples. One is George Perry’s finding, reported elsewhere in this volume, of a much larger market response to unexpected changes in the payroll employment survey, a series that the Federal Reserve is known to favor, than to unexpected changes in the household employment survey, although Perry finds that the latter has comparable information content about the state of the economy. The other is the exaggerated reactions of markets to announcements of the money stock aggregates after 1979, when the Federal Reserve began to emphasize growth of the money stock as an indicator of its monetary stance.

On a more speculative note, the authors report several strands of evidence suggesting that the increase in Federal Reserve transparency in recent years may have been accompanied by a reduction in the information contained in market signals. They see the recent finding by Flint Brayton, John Roberts, and John Williams of a flattening Phillips curve as suggesting that goods prices have deteriorated as an indicator of how close the economy is to capacity. And although they see increasing evidence that

the adoption of inflation targeting is associated with stronger anchoring of inflation expectations, the authors suggest that the flip side of “well anchored” is “uninformative”: prices that are well anchored are also prices that may have low signal value, revealing little about underlying shifts in the economy. They report studies finding that surveys of inflation expectations in inflation targeting countries tend to be unresponsive to a moving average of actual inflation, whereas such surveys report large and statistically significant responses to inflation in countries without formal inflation targets. The authors also find qualitatively similar differences between targeting and nontargeting countries with respect to the forward inflation premium implicit in the difference between nominal and price-indexed bonds. But it is hard to know whether these reductions in either survey or market responses mean there is less information in recent inflation, or simply that future inflation is less related to current experience in targeting countries.

Having described why increased central bank transparency might actually reduce informational efficiency, the authors analyze some formal models to demonstrate conditions under which this could occur. In the simplest model, each of many agents cares how close his or her action is to an unknown economic fundamental. (The authors normalize variables so that it would be optimal to set the action equal to the fundamental if it were known and the agent cared only about it.) But agents also put heavy weight on how close their action is to their expectation of the average action of others (in a manner analogous to the Keynesian beauty contest). If all agents knew the true value of the fundamental and knew that other agents also knew it, they would all set their actions equal to that value, and all would have zero loss. But suppose each agent has only a noisy, private signal of the fundamental, and suppose that all agents know that their private signals are unbiased and uncorrelated across agents. In the absence of a public signal or of sharing of information about their signals, the best any agent can do is to set his or her action to the value of his or her signal—the agent’s own best guess of both the true value of the fundamental and the actions of other agents, all of whom will be using their signals in the same way.

Now suppose the central bank makes public its forecast. This additional information will affect the typical agent’s expectation of the fundamental, making it a weighted average of the public forecast and the private signal. Absent any concern about other agents’ actions, that would also define each

agent's action given the public signal. But since agents care about how far their own actions are from the action of the average agent, they will take account of how other agents will alter their behavior. Furthermore, since each agent expects that other agents will respond in the same way, the response of all agents will give still more weight to the public signal than would be warranted if each simply modified his or her expectations by the direct information contained in the public signal.

The results from this model illustrate the intuition gained from the authors' earlier discussion. An increase in the precision of the public signal unambiguously increases the average distance of actions from the true fundamental. This reflects the fact that the errors in the private signals average out—in the absence of the public signal, those errors might be large, but they would sum to zero. The fact that agents weigh the public signal in deciding their action introduces a common error in every agent's action. The more agents care about whether their actions are close to the average action of others, the greater the weight they will assign to the public signal, and the further the average action will be from the true value of the fundamental.

The authors show that the central features of this simple beauty contest model are present in three economic models in which information plays an important role. In each of these models, two of which are macroeconomic and one financial, agents' actions depend on their expectations of the average actions of others. The most familiar is the Lucas-Phelps "islands" economy, in which islands serve as a metaphor for different regions or sectors of the economy. Agents' supply on each island depends on their expectations about the average of prices across all islands, and demand on each island depends on a fundamental (the money supply) and the price on that island. In the resulting equilibrium the price on an island depends on expectations of the fundamental and the average price across islands in an equation of exactly the same form as the beauty contest. The authors note, however, that in the original version of this model the central bank has full control over the fundamental. Michael Woodford's model of price behavior in a monopolistically competitive world is a second example. In this model firms' expectations of the average prices set by other firms matter because of price competition, and the fundamental is nominal marginal cost.

The third example is a model of financial market prices constructed by the authors with Franklin Allen. Models of financial prices are complicated by the fact that they are forward looking—today's price is the

expected payoff at some future date. Thus the price today is the average expectation today of the average expectation tomorrow, and so on, up to the average expected fundamental. In this context public information has a disproportionate effect, and the law of iterated expectations fails. In particular, as in the simple beauty contest example, the price deviates systematically from the average expectation of the fundamental value, being biased toward the commonly shared information.

In none of the models discussed so far does deterioration in the quality of market information affect the accuracy of the central bank's own information. But, as the authors note, the central bank's information comes from the economy itself and reflects the actions taken by private agents. This suggests the possibility that greater transparency could actually degrade the bank's own information. The authors show this possibility using a model in which the underlying fundamental of the economy follows a random walk, and the central bank's information comes from a noisy signal of the average of private actions in the past. This information about past private actions is assumed to be unavailable to private agents; this simplifies the model but is not crucial to its results. However, each private agent receives a noisy signal of the current value of the fundamental and, as in the beauty contest, cares about both the distance of his or her action from the fundamental and the distance from the average action of other agents. The authors show that the precision of the central bank's estimate of the fundamental depends on the precision of the private signals, their own signal of past private actions, and the variance in the innovations to the fundamental. And in this model it is indeed the case that the central bank has more precise information about the value of the fundamental if it does not disclose its own estimate of that value.

The authors are cautious in drawing strong conclusions about the welfare effects of increasing transparency. It is true that they have shown how, by making more information public, the central bank may actually reduce the accuracy of the information available to private agents and of its own information. But this possibility depends crucially on agents' wanting their own actions to be close to those of others, in effect coordinating their behavior. The authors briefly review three recent papers that explicitly address whether increased coordination of actions increases or decreases welfare, and find that they give varying results. Thus they conclude that the choice of models has a pivotal effect on the conclusion. Perhaps most important, the welfare costs of degrading the information available to the

central bank will depend crucially on how damaging that loss is to the central bank's ability to stabilize the economy, a question not addressed in their models. The authors do not see their paper as resolving whether any particular forecast or type of information should be disclosed. Rather their objective is to point out the trade-offs involved in central bank communication with the public and to draw attention to the two-sided nature of the debate.

SINCE THE MID-1990S the overall performance of the economy has surpassed most expectations. Unemployment fell well below levels most analysts had regarded as the economy's natural rate, inflation has not emerged as a problem, the recent recession was mild, and GDP growth has averaged nearly 4 percent a year. The most striking development throughout this period has been the rapid growth of labor productivity. After averaging 1.46 percent between 1975 and 1995, the annual rise in nonfarm business output per hour worked has averaged 2.83 percent over the past decade. The most disappointing development has been the failure of the median real income of workers to rise by anywhere near as much as productivity. Such a failure can arise either from aggregate labor income failing to keep up with total output, or from increases in the inequality of income that result in the median income rising more slowly than the mean income. In the second article of this issue, Ian Dew-Becker and Robert Gordon analyze both possibilities, examining how labor's share of aggregate income has varied in recent years and how the distribution of income has dramatically changed.

If real compensation per hour grows at the same rate as labor productivity, labor's share of output will be unchanged. Dew-Becker and Gordon show that, over the entire period since 1954, that is about what has happened: economy-wide real compensation per hour rose an average of 2.1 percent a year and GDP per hour an average of 1.9 percent a year. However, over shorter intervals real compensation and output per hour have not moved so closely, and labor's share of income has varied more, generally rising in the initial postwar decades and declining somewhat since the early 1980s, but with substantial shorter-run fluctuations.

Changes in labor's share can also be inferred from the overall rise in prices relative to total unit labor cost. Hence the effects of changes in productivity growth on labor's share are implicit in their effects in price and unit labor cost equations. This leads the authors to study the impact of

productivity on income shares by introducing variations in productivity growth into Gordon's previous model of inflation dynamics. This model, a recent version of which was published in a 1998 Brookings Paper, already included a number of explanatory variables: the gap between the actual and the natural rate of unemployment (where the latter is free to vary over time); long lags on the dependent variable; a number of supply-shock variables; and price control dummy variables for the Nixon years. To this list the authors add the acceleration in labor productivity, lagged one quarter and five quarters, and a supply-shock term for medical costs. They use this revised model to explain changes in the personal consumption expenditure deflator over 1962–2005 and find that the productivity terms are jointly significant, with faster productivity growth decreasing inflation, and that increases in medical costs cause significant price increases. They also show that the equation performs better in out-of-sample simulations when these two terms are included.

The authors next estimate a form of their equation that includes lagged changes in the labor share of income as an additional explanatory variable. They use this form in regressions explaining, alternatively, inflation in the personal consumption and nonfarm business deflators and inflation in trend unit labor costs. The equations for the change in unit labor cost are less satisfactory than those for price inflation, with the coefficients on productivity acceleration still negative but not significant. However, the differences in coefficient estimates between the unit labor cost and price equations do imply some interesting dynamics. An acceleration in productivity will initially lower inflation more than it reduces the growth in unit labor costs, thus adding to the change in labor's share. This effect vanishes after a year. Using a simulation of the nonfarm business deflator and unit labor cost equations in which they suppress the productivity deceleration of 1965–80, they show that the deceleration of productivity added 1.3 percentage points to the average inflation rate over the period and 2.7 percentage points by the end of the period. The effect on trend unit labor costs was slightly larger, contributing 0.18 percentage point to the average annual change in labor's share. The same simulations for 1995–2000 show that the acceleration of productivity in that period lowered the change in labor's share by 0.19 percentage point a year. It reduced the annual rate of inflation by an average of 1.2 percentage points a year over the entire period and by 1.7 percentage points a year by the end of the period. The authors conclude that this acceleration of productivity was

the single most important factor permitting the Federal Reserve to maintain low interest rates in this period.

In the second part of their paper, the authors turn to Internal Revenue Service (IRS) data from tax returns to examine the experience of the median worker and the distribution of income more generally. Although labor's share of income changed little over the decade from mid-1995 to mid-2005, the period of surprising productivity growth, the authors show that the experience of the median worker was sharply worse than growth in average compensation, which determines labor's share.

Most studies of income inequality are based on data from the Current Population Survey (CPS). Although the CPS data have several advantages over the IRS data for incomes up to the top decile, its income categories provide no information about the distribution of incomes within this decile. The IRS data, in contrast, identify incomes without an upper limit and come from an annual sample of over 100,000 returns, with oversampling of returns at the very top of the distribution. They thus permit a far more revealing analysis of distribution within the top decile, which turns out to be important for understanding recent developments.

The authors find that the labor income distribution in their IRS data shows the same broad trends over time as the distributions previously reported by David Autor, Lawrence Katz, and Melissa Kearney, who used two alternative samples of the CPS data. The main trends can be characterized by the ratios of the 50th to the 10th percentile (the 50/10 ratio) and the 90th to the 10th percentile. In all measures the 50/10 ratios were relatively constant throughout the 1966–2001 period. By contrast, in all measures the 90/10 ratios rose substantially over this period, indicating that all of the change in inequality over the period occurred above the 50th percentile. The bottom half of the distribution received only 12 percent of the gain from 1966 to 2001.

Using the IRS data, the authors are able to refine these estimates, tracking distribution changes within the top decile where the income growth was most concentrated. The results are striking. Over the period 1966–2001, they find that, whereas the 90/10 ratio rose by only 34 percent, the 99/10 ratio rose by 87 percent and the 99.9/10 ratio by 185 percent, indicating an increase in the concentration of income at ever higher quantiles. More of the aggregate income increase went to the top 1 percent than to the bottom 50 percent, and more went to the top 0.01 percent than to the bottom 20 percent. Between 1997 and 2001, the period that includes the recent produc-

tivity acceleration, the changes are generally similar to those observed over the entire sample period, except that the gains are even more concentrated at the very top: 24 percent of the gains went to the top 1 percent, compared with 13 percent to the bottom 50 percent, and twice as much went to the top 0.01 percent as to the bottom 20 percent.

The authors estimate the increases in real average hourly compensation in each income decile and compare them with labor productivity growth. Over the entire 1966–2001 period, productivity growth averaged 1.57 percent a year. No decile below the 90th percentile enjoyed growth in real compensation as fast as that. But the shortfalls were greatest in the years between 1972 and 1997. At both ends of their sample period, 1966–72 and 1997–2001, real hourly compensation exceeded productivity growth at each income decile. For the median worker, over the entire sample period real compensation growth averaged only 0.76 percent a year, less than half the growth rate of productivity. And for the twenty-five years from 1972 to 1997, real compensation of the median worker declined, falling behind productivity growth by about 1.5 percentage points a year.

The authors conclude by examining possible causes of the changing income distribution, informing their discussion with insights from their analysis of IRS data. They first review the explanations of previous analysts. Skill-biased technical change, usually associated with the rise of computer technology and related applications, has been commonly offered as an explanation for widening income disparities. However, David Card and John DiNardo, along with several subsequent authors, have rejected this explanation, pointing out that the timing seems wrong: neither the development of computer technology nor the slowdown in productivity correlates in time with the rise in inequality. Card and DiNardo's favored explanation is the sharp drop in the real minimum wage in the early 1980s, which does correspond closely in time with the sharp rise in their measure of inequality. Autor, Katz, and Kearney concluded that demand for labor has increasingly favored tasks associated with occupations in the top income deciles. Claudia Goldin and Robert Margo have stressed the need to take a long view; they explain the flattening of the income distribution during 1930–70 as well as the growing inequality since then, identifying three main developments that fit both: the rise and fall of unionization, the decline and recovery of immigration, and the decline and recovery in the share of imports.

Dew-Becker and Gordon agree with the critiques of the skill-biased

technical change hypothesis, and they acknowledge the likely role of some of the alternative explanations. But they note that none of the explanations accounts for the extraordinary increases they have identified in the top decile of the distribution. Their own explanation attributes these changes to two developments that might be expected to generate exceptional income gains at the very top. The first is the “superstar” phenomenon, first identified by Sherwin Rosen as an equilibrium phenomenon in unusual markets. The second is the explosion of incomes of CEOs and other top corporate officers.

Superstars in the entertainment world earn huge incomes because of a confluence of factors operating in the contemporary economy. One is that consumer preferences are highly skewed—as Rosen put it, “hearing a succession of mediocre singers does not add up to a single outstanding performance.” Another is that a series of technological innovations, from phonographs and radios to movies and television, have magnified the potential audience and thus the revenue that can be earned by the few at the very top. Both these factors have increased in importance in recent decades. But are there enough superstars to matter in the aggregate? The authors piece together evidence suggesting that there are. In 2001 the 99.99th percentile received \$83 billion in income. *Forbes* magazine’s list of “The Celebrity 100” accounts for \$3.1 billion of that, and the total payrolls of professional football, baseball, and basketball teams account for another \$7.0 billion. When one takes into account that earnings are also highly skewed at the top in other fields, such as law and finance (although data for these are not readily available), it is reasonable to conclude that superstars can account for a meaningful part of aggregate income at the very top of the distribution.

The authors see the growing pay premiums of CEOs as a quite different phenomenon—one not easily seen as a conventional market outcome. They cite one study showing that the ratio of average CEO pay to average worker pay rose from 27 to 300 between 1973 and 2000 before falling to 237 in 2001 when the stock market crashed. Other studies also show substantial although less spectacular increases. For more recent years they cite a study by Lucian Bebchuk and Yaniv Grinstein showing that the compensation of the top five executives in a sample of firms rose from 5 percent, on average, of the firm’s total earnings to 13 percent between 1993–95 and 2000–02, an interval of substantial growth in corporate earnings. Dew-Becker and Gordon show that the salaries of top executives constitute a substantial share of income at the very top of the distribution. The average pay of the

top five executives in the 1,500 firms of the Standard and Poor's 500, the Mid Cap 400, and the Small Cap 600 indexes was \$32 million, for an average of \$6.4 million each. This adds up to \$48 billion, more than half of the income in the 99.99th percentile.

Dew-Becker and Gordon conclude that the superstar and CEO hypotheses together explain much of the surge at the very top of the income distribution. Furthermore, the importance of the very top has grown and now accounts for a substantial part of the recent discrepancy between the growth of average labor compensation and the disappointing growth in compensation of the median worker. Looking ahead, the continued global spread of entertainment technology and globalization of financial and commercial markets could well exacerbate these trends.

THE SUDDEN AND DRAMATIC collapse of the energy trading company Enron Corporation in the fall of 2001, ending with its bankruptcy on December 2 of that year, resulted in thousands of Enron employees losing their jobs and, simultaneously, a substantial fraction of their retirement assets. At the end of 2000, current and former Enron employees held 62 percent of the savings in their 401(k) retirement plans in Enron stock, which became worthless less than a year later. A similar fate has befallen employees of WorldCom, Global Crossing, Polaroid, Kmart, Lucent, and Providian, among others. Unlike defined-benefit pension plans, 401(k) plans and other defined-contribution plans currently face no restrictions on how much of an employee's account may be invested in the stock of the employer. In response to widespread public concern following these well-publicized debacles, many bills have been introduced in Congress to regulate employer stock holdings in 401(k) plans. But rather than place outright limits on such holdings, most proposals would expand employees' rights to sell employer stock in their account or require the employer to inform them of the risks of concentrating their holdings. In the third paper of this issue, James Choi, David Laibson, and Brigitte Madrian examine the likely effects of such "empowerment and education" provisions. Their work focuses on several recent natural experiments that provide evidence about how employees utilize investment choice when given the opportunity, and analyzes how employees respond to demonstrations of the risk of placing all their retirement eggs in one basket.

The natural experiments involve five large publicly held companies where restrictions on the portfolio choices of employees were relaxed,

either because employees crossed age or tenure thresholds or because their company changed its rules, allowing them to diversify. The data for these episodes come from Hewitt Associates, a large benefits administration and consulting firm, and include a series of cross sections covering all persons employed by the firm at the time restrictions were relaxed; the information available for each employee includes stock holdings and demographic data such as birthdate, hire date, and compensation.

As regulated under ERISA (the Employee Retirement Income Security Act of 1974), 401(k) plans may differ across companies in the restrictions on ownership of company stock and in the options that must be available to employees. If the plan is not combined with an employee stock ownership plan (ESOP), the company may not require that more than 10 percent of the employee's own contributions be invested in the employer's stock. If the plan is combined with an ESOP, this restriction does not apply. Neither type of plan limits how much of the employer's matching contributions may be directed into employer stock. However, all assets within an ESOP are subject to explicit diversification options. Employees with ten years of tenure must be allowed to diversify at least 25 percent of their employer stock holdings once they reach age fifty-five and 50 percent after reaching sixty.

Only one of the five firms studied by the authors (the firm they designate as company D) allowed just this minimum flexibility during the study period, and it did so only until early 2002, when it eliminated all such restrictions. Strikingly, however, before the minimum requirement was lifted, the average company D employee who could have diversified never held less than about 90 percent of his or her total portfolio in company stock. When holdings are weighted by the size of the portfolio (that is, simply dividing total employer stock holdings by total assets), the results are little different, suggesting that even those employees with longer tenure or higher salaries sought little more diversification than the average. A very similar pattern emerges for companies B and C, both of which allowed greater flexibility than the law required: both simple-average and dollar-weighted holdings of company stock by their employees seldom dipped below 90 percent of the total.

The authors also examine investment behavior following changes in restrictions by companies A, D, and E. Company A voluntarily changed its rules in 2002, allowing employees to diversify up to 25 percent of the employer stock in their match account at age forty-five, 50 percent at age

fifty, and 100 percent at age fifty-five. The proportion held in company stock was typically lower, the older the employee. But, on average, employees of any age never sold more than a third of the company stock they were permitted to sell. Dollar-weighted averages fell somewhat more, indicating that wealthier employees were more eager to diversify. Company D, as noted above, eliminated all diversification restrictions in early 2002, but it continued to direct the employer contribution to company stock; maintaining a diversified portfolio thus required the employee to trade out of company stock on an ongoing basis. Company E adopted the same policy in mid-2002. The authors show that, at company D, the fraction of all employee 401(k) balances held in company stock drifted steadily downward. Yet, even after two years, 84 percent of holdings were still in company stock. Data for company E were available only for the end of 2003, but the picture is similar: the average participant under age fifty still had 90 percent of matching balances in company stock, and that share still exceeded 80 percent among older workers. The authors see these natural experiments as painting a consistent picture: relaxing diversification restrictions does decrease employer stock holding, but only modestly, especially among younger employees.

If the failure of employees to diversify their portfolios reflects a lack of awareness of the dangers of placing such a large wager on their own firm's success, perhaps better information about those dangers is what is needed. Six out of twenty-one 401(k) reform bills recently introduced in Congress (listed in an appendix to the paper) take this educational approach: these bills would require firms to inform their employees when their holdings of company stock exceeded some percentage or, in some cases, provide information about the desirability of diversification or the risks of concentrating their assets in company stock. But would such provisions make any difference? Previous studies have shown that financial education alone generates only small changes in 401(k) outcomes. None, however, have investigated the effect of education by (vicarious) experience, that is, of witnessing the consequences of others' "financial follies." This leads the authors to examine how media coverage of the financial debacles of Enron, WorldCom, and Global Crossing affected the portfolio choices of employees at other companies.

The authors painstakingly construct four different proxies for daily media coverage of the financial collapses of these three companies, beginning on October 16, 2001, the day Enron's market capitalization fell by

\$1.2 billion. After culling for duplications, the authors identify over 12,000 relevant newspaper stories and 1,900 television stories between October 16, 2001, and December 31, 2003. Counting only those stories that include the term “401(k)” reduces these numbers to 761 and 144, respectively. Recognizing that the impact of one day’s stories may last several days, the authors construct for each of these four story counts a “news impact” variable that assumes that the impact decays by a constant percentage each day after publication or broadcast. News coverage varied widely over the period, with the most active coverage occurring from mid-December 2001 through the end of February 2002, when Enron and Global Crossing filed for bankruptcy, and again in June and July 2002, when WorldCom disclosed major accounting irregularities and shortly thereafter filed for bankruptcy.

How did this news affect stock holdings in 401(k)s? For a sample of firms whose employees held a total of approximately \$45 billion in 401(k) assets, the authors show that the share of those assets held in employer stock began a declining trend around the time of Enron’s collapse. However, as the authors note, the overall stock market was also declining around the same time, and indeed the share of *equities* held in company stock showed much less of a downward trend. The authors find it striking that this share barely moved during the time that the Enron scandal was unfolding.

The authors next investigate whether employees’ net transfers are related to the news impact variable. They first look at daily data on the value of aggregate net transfers of employer stock (excluding payroll contributions in company stock) as a ratio to the preceding day’s aggregate employer stock balance. A time-series graph of this variable indicates that the volatility of net trading increased following Enron’s collapse, but the mean direction of trades did not obviously shift in either direction. The authors run separate regressions of this series on each of the four news impact variables, including as additional variables current and lagged market returns on the Standard and Poor’s 500 stock index, day-of-the-week dummies, and a polynomial time trend. They find that the news variable has a significant and correctly signed coefficient in all the regressions, but the coefficients are small. Adding up the effect of all news after Enron’s collapse, they find that none of the specifications imply more than a 2.4-percentage-point drop in employer stock holdings as a share of the total portfolio. The authors recognize that some of these employer stock holdings are constrained, but even when calculated as a percentage of completely unre-

stricted accounts, the reductions in company stock holdings are less than 6.2 percentage points.

To further investigate the importance of holding requirements, the authors study two large companies for which they have daily trading data. For one of these, company D, which until 2002 set restrictions on diversification out of the employer match account, they examine only the balances in the unrestricted employee accounts. The other company, company F, has no diversification restrictions anywhere in its 401(k) plan, and indeed matching company contributions mirror whatever employees choose for their own contributions. The analysis of data for these firms provides even less support than the aggregate data for the hypothesis that employees respond to bad news about other companies by reducing their investments in company stock. Both the fraction of account balances and the fraction of employee contributions in company stock trended upward during the period and show little evidence of response to news. In equations testing for a response, the regression coefficients are small, sometimes of the wrong sign, and never significant.

The authors also investigate whether the Enron scandal was more salient for workers living in Texas than elsewhere. More than a quarter of Enron's employees worked in Houston, and the company cut more than 4,000 jobs at its Houston operations in 2002. Two large firms in the authors' sample have a substantial number of employees both in and outside of Texas, making it possible to examine what difference proximity makes to the behavior of employees of the same company. It seems to make none: comparing employees with the same date of plan enrollment, the authors find no significant differences in the portfolio allocations of Texan and non-Texan employees (or, in the case of one firm, employees working in or outside of Houston).

The authors conclude that empowering employees to trade out of employer stock and educating them about the risk of employer stock will have only a small effect on the concentration of employer stock in 401(k) holdings. They note that this conclusion is consistent with earlier studies showing that consumers are often passive and that educational interventions yield remarkably small changes in behavior. The authors believe their results cast doubt on the value of empowerment and education in proposals for 401(k) reform. Most economists and financial advisers believe in the advantages of diversification and counsel that stock in one's own company should make up only a modest fraction of the typical retire-

ment account. The authors grant that economists generally dislike paternalism and coercive policies. But they note that specifying a default option would place no constraint on freedom of choice—the employee can always override the default. The default could be a well-diversified asset allocation, or it could simply be required to meet the same fiduciary standards that apply to defined-benefit pension plans. The results both of this paper and of earlier work suggest that many employees would stick with the default, and in doing so would end up with much sounder portfolios than the typical 401(k) participant holds today.

ENERGY IS ESSENTIAL TO modern economies, and as China, India, and other poorer economies develop, world demand for energy will rise rapidly. Growth in demand has already been one important factor driving up the price of oil and natural gas during 2005. High prices, in turn, have encouraged the expansion of reserves in existing fields, the search for new fields, and investment in unconventional sources of petroleum such as tar sands and oil shales. Absent major political disruptions in supply, many observers believe the world's growing energy demand can be met from these sources and other existing fossil-fuel-based technologies for years or decades, at costs that are within the range of historical experience. However, there is increasing consensus that satisfying growing demand in this way risks severe environmental damage, and that avoiding this risk will require policy interventions and more sweeping changes in how energy is produced and used. In the fourth article of this volume, Klaus Lackner and Jeffrey Sachs consider these longer-term prospects and the uncertainties surrounding them, analyze the potential for protecting the global environment through existing technologies and foreseeable new ones, and discuss what the broad aims of policy for the twenty-first century should be.

Lackner and Sachs start with a technical overview of energy use and its environmental impact. Fossil fuels—oil, natural gas, and coal—provide 85 percent of current U.S. and world energy use. Hydroelectric and nuclear power provide almost all the rest. Almost all the fuel used for transportation today comes from oil, but crude oil reserves need not limit transportation in the future. Known technologies permit the substitution of one fossil fuel for another. Oil shales and tar sands are already being used to produce petroleum. And a method known as the Fischer-Tropsch process has already been used for several decades to convert coal into gasoline and diesel at costs near current oil prices. Thus, although fossil fuels are, in principle,

an exhaustible resource, known coal reserves are so vast that they could satisfy the world's transportation needs well beyond the present century. Biomass can also be used to produce fuel for transportation, but the authors see its quantitative potential as limited.

Although the authors thus see no relevant limit, at least in this century, to economic development from a possible decline of global oil and gas resources, they do perceive a pressing environmental imperative to reduce the world's reliance on fossil fuels. Combustion of fossil fuels produces carbon dioxide (CO₂), the principal greenhouse gas, as a byproduct. The capacity of the earth to absorb this CO₂ is limited, although how limited is open to debate. The problem is more complicated than public discussions of "global warming" often suggest. As the authors explain, besides the greenhouse effect, increasing atmospheric CO₂ concentrations will affect various chemical, climatological, and biological processes on a vast scale. Some likely consequences include major changes in winds, rainfall, and ocean currents; rising ocean levels; the possible destruction of ocean ecosystems; the extinction of species with limited mobility; the enhanced transmission of diseases, such as malaria, that are now constrained by climate patterns; dislocations from changes in agricultural productivity; and an increase in extreme weather events. They also warn that such developments may not be gradual, because feedbacks from relatively small increases in CO₂ concentrations can be large. For example, if rising temperatures dislodge the ice sheets of Antarctica or Greenland, the paths of major ocean currents could be affected; melting permafrost could release buried methane, itself a greenhouse gas, thus adding to the total concentration of greenhouse gases in the atmosphere.

Uncertainties about the size and timing of these and other effects of increased CO₂ are large. Existing climatological models cannot adequately explain what has already happened, let alone agree on quantitative projections under alternative paths of CO₂ emissions. In 1992, when these uncertainties were even greater than today, the United Nations Framework Convention on Climate Change committed all signatories, including the United States, to stabilize "greenhouse gas concentrations . . . at a level that would prevent dangerous anthropogenic interference with the climate system." But the United States and some other countries have since rejected this approach. The Bush administration in particular has argued that the costs of mitigation would exceed the benefits and has rejected any specific targets.

Lackner and Sachs believe that current evidence supports setting a CO₂ target despite the remaining quantitative uncertainties, and they summarize what current knowledge suggests about reasonable target levels. There is now widespread agreement that the continued buildup of CO₂ poses risks to the global ecosystem and to human society. The preindustrial concentration of CO₂ in the atmosphere was 280 parts per million (ppm), and this has already risen to 380 ppm. They report near unanimity in the scientific climate community that a tripling of the preindustrial concentration, to 840 ppm, would pose catastrophic risks and that global economic growth over the next century is likely to raise concentrations to that level or even higher if unfettered energy market forces are allowed to determine future production and use. In light of this danger, the authors accept the consensus among European governments and analysts to limit the atmospheric concentration of CO₂ to between 450 and 560 ppm. (The latter is referred to as the 2× standard since it is twice the preindustrial concentration.) Along with setting a target, the authors reason that a robust policy should be global, since global emissions are what matter, and should begin promptly so as to minimize the present value of costs.

Carbon emissions will inevitably continue to rise for some time, as will the CO₂ buildup in the atmosphere. But the net carbon buildup must be reduced in the future and eventually brought down to zero if any reasonable CO₂ target level is to be maintained. The authors describe three main technologies that have the potential eventually to satisfy world energy needs with zero carbon buildup. The first is carbon capture and sequestration (CCS), in which the use of fossil fuels continues but the resulting carbon emissions are captured and stored rather than released into the atmosphere. Developing and applying CCS is the most immediate need, because it will allow continued use of the existing energy infrastructure. The other two technologies are nuclear and solar energy, both of which must overcome substantial drawbacks if their role in meeting future energy needs is to expand.

Nuclear power today provides about 18 percent of the world's electricity, but the authors see current estimates of global uranium resources as too small to support a greatly expanded role for conventional nuclear technology. If it can be perfected, breeder reactor technology would greatly relax these resource limitations, but it would also increase the risk of nuclear weapons proliferation, since the breeder process produces weapons-grade material. Fusion energy would pose much less risk, but it

remains only a theoretical possibility. Solar energy is, in principle, virtually unlimited but far too expensive with current technology. The authors regard it as a potentially important source of energy, but only for the very distant future.

The authors report that CCS is already applied on a small scale in some industries, and they discuss the costs and improvements in techniques needed to apply it on the massive scale required to achieve net zero carbon growth worldwide: over a century, permanent storage would be needed for thousands of gigatons of CO₂. Options include disposal in the ocean, storage underground, and chemical fixation as a solid. The authors discuss the pros and cons of each, including the experience to date, on a small scale, with applications such as injection of CO₂ in enhanced oil recovery. They are skeptical about disposal in the ocean because it would not be truly permanent and would risk disruptive environmental changes to the ocean itself. They conclude that carbon storage should start today with underground injection and eventually move to chemical processing of CO₂ into stable solid carbonates. Injection is economically attractive today, but the cost of chemical processing would have to fall by a factor of four or five to keep the cost of energy within 30 percent of today's cost.

Capturing carbon is most economical in large fossil-fuel-using applications such as power plants. When designed into a new plant, it entails roughly a 30 percent efficiency penalty. The authors also discuss potential alternative designs that would substantially alter how power plants produce their energy from coal. They note that carbon capture for dispersed or mobile sources of CO₂ is not a realistic option, but that there is substantial room for reducing emissions from such sources, for example from the adoption of low-emission vehicles and the substitution of electricity for oil or natural gas for home heating.

The authors next turn to scenarios of future world economic growth, to show how the production and use of energy are likely to evolve if present patterns continue, and what the implications are for the buildup of CO₂ in the atmosphere. They then show how changes in production and use could limit that buildup. Their scenarios are calculated from a disaggregated model of global energy use, which divides the world into eight regions: the United States, Western Europe, other developed economies, the transition economies, China, India, other emerging Asian economies, and all other emerging economies. Each economy uses primary energy—oil, gas, coal, nuclear, and renewable—directly and indirectly through the use

of electricity produced from primary energy. U.S. income per capita is projected to grow by 1.7 percent a year, its long-term trend growth rate, and incomes per capita in other countries are projected to converge toward the U.S. level throughout the projection period.

In the baseline, “business as usual” scenario, primary energy demand in each region grows in proportion to income per capita less an efficiency saving of 1.5 percent a year. In each country primary energy use and electricity use are assumed to grow in proportion to overall energy use. World GDP measured in 2002 dollars grows from \$46 trillion in 2002 to \$278 trillion in 2050 and \$910 trillion in 2100. The share of today’s developed regions in world GDP falls sharply, from 60 percent in 2002 to 29 percent in 2050 and 21 percent in 2100. Although by 2050 efficiency gains cut energy use per dollar of GDP by 50 percent, world demand for primary energy grows 2.8-fold by 2050 and 4.3-fold by 2100.

The authors estimate the effect on the global climate of this rise in energy use, taking account of the fact that part of each year’s carbon emissions will remain in the atmosphere, part will be dissolved in the ocean, and part will be incorporated in living plants and soil, with the latter two effects taking place gradually. They calculate that fossil-fuel-based emissions will rise from the current level of 5.8 gigatons of carbon (GtC) a year to 26 GtC in 2100, raising the concentration of atmospheric carbon to 554 ppm by 2050 and to 886 ppm by 2100. They note that these projections are highly conservative, for at least two reasons. First, declining petroleum reserves will likely lead to substitution of coal and other fossil fuels; this would raise carbon emissions above the baseline because these fuels emit more carbon per unit of energy than does petroleum. Second, if the developing world’s use of automobiles grows faster than its GDP, as it almost certainly will, its energy demand will rise much faster than in the baseline projection.

Lackner and Sachs’ alternative scenario considers what they see as the best way to avoid the unacceptable business-as-usual outcomes. For the immediate future, it involves two main initiatives that are immediately available: sharply improving efficiency in transportation by embracing hybrid technology in automobiles, and introducing CCS technologies in facilities where coal is used on a large scale. Although the authors project that hybrid vehicles could double fuel efficiency—a much greater improvement than present hybrid models can provide—making their adoption cost effective to consumers, they nonetheless argue for subsidizing hybrid auto-

mobiles in order to achieve their full potential to reduce carbon emissions. In their alternative scenario, and with their projected gains in fuel efficiency, converting entirely to hybrid cars would reduce CO₂ concentrations in 2050 by 4 percent from the baseline. As for CCS, although its application is in its infancy, the technologies themselves are known and ready to be exploited. The authors calculate that the total costs, including capital costs, of CCS in power generation are in the range of 1 to 3 cents per kilowatt-hour. In their alternative scenario, by 2050 the full implementation of CCS would cost somewhere between 0.1 and 0.3 percent of world GDP but would reduce CO₂ emissions by an estimated 17 Gt a year. If CCS were phased in starting in 2006, the atmospheric concentration of CO₂ in 2050 would be reduced by about 8 percent. The authors emphasize that initiatives in both hybrid transportation technology and CCS should be taken without delay. Waiting allows emissions to continue rising, invites continued construction of long-lived non-CCS plants, and postpones the learning and cost reductions that will surely come once better technologies are actively pursued.

Although these first steps would slow the trajectory of CO₂ buildup, they would not keep it from eventually rising to dangerous levels. For the second half of the century, therefore, further improvements will be needed. These will have to come from reducing point-source emissions—those from cars, homes, or businesses that are too small to use CCS. Much more widespread electrification or the development of other noncarbon energy carriers, such as hydrogen, could eliminate point-source emissions if technology can make them cost effective over this more distant horizon.

The authors emphasize that growth in the developing world will be the main driver in carbon buildup. In their base case, the developing countries account for almost 60 percent of emissions in 2025 and 70 percent in 2050. China and India are the countries where the greatest increase in energy consumption will take place and therefore where most new power generating facilities will be built. And it is much less expensive to reduce emissions by building new facilities with the new technology than by retrofitting old ones. Thus any successful program to reduce carbon emissions will have to center on these and other countries in the developing world. The authors note that either a uniform tax on carbon emissions in all regions or a system of global tradable permits—both of which have been proposed—could, in principle, provide the needed incentives to restrain emissions. However, they recognize that negotiations to do either may be difficult, and they

believe that the only practical way to reduce emissions may be to secure the cooperation of the limited number of decisionmakers who license new power plants and set new efficiency standards. Whatever route is taken to achieve global cooperation, since most of the carbon reduction will have to be done in regions that can least afford the added cost, the authors stress that part of the incremental cost of CCS in China, India, and other developing nations should be borne by the high-income countries.

THE MONTHLY EMPLOYMENT REPORT of the Bureau of Labor Statistics (BLS) is the government's most widely anticipated statistical report. It includes the most current information available on labor market developments from two sources, a survey of households and a survey of payrolls, and so offers the earliest and broadest picture of how the overall economy is performing. The household survey, based on monthly responses from about 60,000 households, measures the total number of people employed and unemployed and provides a wealth of more detailed information, much of it disaggregated by demographic characteristics. The payroll survey, based on monthly reports from about 400,000 business establishments (and revised annually on the basis of more complete information), measures the total number of nonfarm jobs, hours worked, and average hourly earnings, all disaggregated by industry. Each survey thus provides a measure of overall monthly employment growth. The two measures often differ by a substantial amount, and the measure from the payroll data is widely regarded as the more reliable. It is featured on the front page of the BLS's monthly release and in press and television coverage, and Federal Reserve Chairman Alan Greenspan has publicly stated his preference for it. In the last paper in this volume, George Perry questions whether this overwhelming preference for the payroll estimate of aggregate employment is warranted and concludes it is not.

Perry starts by looking for evidence that policymakers and markets in fact respond to the payroll data and ignore the household data. He first analyzes how Treasury yields respond to the two measures of monthly employment change, using regressions of changes in market yields of Treasury securities from just before to shortly (either twenty-five minutes or twenty-four hours) after the employment release on the surprise in the employment change from each survey. The "surprise" is taken as the difference between the actual change and the expected change as reported in surveys of analysts; separate regressions are performed for Treasuries with

different maturities, ranging from three months to two years. The payroll data are clearly favored: they have substantial explanatory power in most specifications, including those that use the household data and changes in unemployment as additional explanatory variables. The household data, in contrast, explain little of the variation in yields, even when they are the only explanatory variable, and they contribute nothing when entered together with the payroll data.

To examine which employment measure monetary policymakers respond to, Perry performs regressions that explain changes in the federal funds rate, the Federal Reserve's principal policy instrument, by changes in the two employment measures between meetings of the Federal Open Market Committee. In all specifications—those using each employment measure alone, both measures together, and both measures together with the inflation rate and changes in the unemployment rate—the payroll data again dominate not only the household data but all the other variables as well. In regressions that do not include the unemployment rate, the household employment data are significant, although much less important than the payroll data.

Having established that markets and policymakers indeed focus overwhelmingly on the payroll rather than the household employment data as an indicator of how the economy is performing, Perry turns to whether this focus is justified. There are a number of definitional differences between the two series, the main ones being the omission of agricultural workers and the self-employed from the payroll data and the counting of multiple job holders only once in the household data. He notes that the estimated sampling error of the payroll data is much smaller but that other sources of error may not be. Reviewing a number of nonsampling problems inherent in the two series, he reasons that the biggest problem may arise in the payroll series' need to account for the birth and death of business establishments. Over 300,000 new establishments are formed each quarter, and accounting for them and allocating their employment, as well as employment losses in establishments that close, poses an especially difficult problem for estimating monthly changes in payroll employment. Perry concludes that the known characteristics of the two series give no decisive reason for considering one more reliable than the other, and he therefore turns to statistical tests. For this purpose he first compares the two series for the period beginning in 1994; the methodology used in the household survey was substantially improved in that year, presumably making this period the most relevant for assessing the series' usefulness now and in the future.

Since policymakers, market participants, and forecasters are interested in the cyclical performance of the economy, Perry first examines how well each measure of employment change tracks contemporaneous changes in output and unemployment, the other two main cyclical variables. He notes that other potential candidates for cyclical tracking tests, such as industrial production or help wanted advertising, are themselves partly derived from the payroll data or refer only to payroll jobs, making them useless for comparing the two employment series. He reasons that, although unemployment is measured from the household survey, regression bias from response error should be minimal since the main response error is between the categories “unemployed” and “out of the labor force.”

Data on GDP and private nonfarm business output, the main measures of aggregate performance, are available quarterly, and so Perry’s first tests use quarterly averages of employment change. The household data explain substantially more of the variation in either output measure than do the payroll data as initially reported. They also outperform the payroll data after the annual revision of the latter, although by less. When entered together with the revised payroll data, the household series has the larger coefficient. Perry also tries an alternative “research” series, constructed by the BLS, that adjusts the household data to payroll concepts. The household series outperforms this measure, too, suggesting that fluctuations in self-employment are important in tracking output. For the regressions tracking changes in unemployment, monthly data are available. Again the household data outperform the payroll data as originally reported, and they do about as well as the payroll data after their annual revision. Quarterly unemployment regressions give parallel results. As in the regressions tracking output, the household data also outperform the research series.

As a further test, Perry estimates simple aggregate employment functions in which current and lagged changes in output, together with deviations in average hours worked from their trend, are used to explain changes in payroll and household employment. In these regressions the household data are better explained when only current and once-lagged changes in output are used, whereas the payroll data are better explained when two additional lags of output are used.

Taken collectively, these results for the period since 1994 provide no support for the widespread belief that the household data should be disregarded in favor of the payroll data as a measure of quarterly or monthly changes in aggregate employment. Repeating the regressions for each of

the four decades before 1994 yields some interesting differences, however. For the two earliest intervals, the decades beginning in 1954 and 1964, the payroll data clearly dominate in all tests. But for the two more recent decades, beginning in 1974 and 1984, the results are mixed, with some specifications and periods favoring one and some the other. Perry notes that the dominance of the payroll series for the earliest decades is consistent with the idea that users came to prefer it then because it correlated best with other indicators of economic performance and never seriously reexamined that choice, influenced perhaps by the household series' larger statistical standard error and its larger monthly volatility.

The empirical evidence for recent decades suggests that averaging monthly changes in the payroll and household data may be the optimal strategy. Perry shows that the monthly volatility of this averaged series is as low as the volatility of the payroll data alone and substantially lower than that of the household data. He examines the performance of the averaged data in tracking equations for the period since 1994. In tracking output, the average formed with payroll data as initially reported does nearly as well as the household data alone, and better than the payroll data alone. The average formed with annually revised payroll data outperforms even the household data alone. The improvements are even stronger in equations tracking unemployment. Here the average employment change, using either revised or unrevised payroll data, dominates either the household or the payroll data alone in both monthly and quarterly regressions.

Perry concludes that users should average the reported changes from the two data sources to get the most useful characterization of monthly and quarterly employment growth. He acknowledges that the BLS cannot report this average as its official estimate of employment growth, because of the definitional differences in what the two series measure. However, he suggests that the BLS give the two measures of employment change equal prominence in its monthly press release.

