The Region
September 2005

Class Action [In Brief]

Economic inequality has increased in the United States—this much is clear. But the reasons, the consequences and the future are far from obvious. Recent research provides some intriguing suggestions.

Douglas Clement
Editor

Thirty years ago, in *Equity and Efficiency: The Big Tradeoff*, economist Arthur Okun argued that societies could not redistribute wealth without a loss in efficiency. A nation might choose to make such sacrifices, he said, but it should be aware of the cost. And in making this argument, he used an especially effective metaphor: a leaky bucket.

Redistribution from one person to another is like carrying water in a leaky bucket, Okun explained: The transfer itself involves waste. Inefficiencies lie in the administrative costs of running the redistribution program and also in the incentive costs—rich people working less because they can't keep the full fruits of their labor and poor people avoiding work because they might lose benefits if they have jobs. Economists ever since have employed the leaky bucket in describing the “efficiency costs” of redistribution policies.

As evocative as Okun's bucket was, another metaphor gained greater prominence (at least among non-economists) in late 20th century debates over income and wealth distribution. In *A Theory of Justice*, published in 1971, Harvard philosopher John Rawls invoked the “veil of ignorance” as the standard of a just society. Decisions about redistribution can only be made objectively, Rawls argued, if we engage in a thought experiment: What allocation would you agree to if you stood behind a hypothetical veil, in a position where you don't yet know what status you'll hold in society—whether male or female, rich or poor, from Manhattan or Maputo?

In this “original position,” he said, our concern would be for the least well-off in society—since we ourselves might end up in that position. A just society is one whose social contract and wealth allocation we would agree to before knowing our position in it, and one in which the welfare of the least fortunate should be maximized.

Liberals embraced Rawls' veil as the foundation for expansion of government poverty programs. But his counterpart at Harvard, philosopher Robert Nozick, would have none of it. Government had no role in such redistributive decisions, he argued in *Anarchy, State and Utopia*. “We are not in the position of children who have been given portions of pie by someone who now makes last minute adjustments to rectify careless cutting,” he wrote. Nozick's libertarian view supported equality of opportunity, not outcomes.

In the decades since these viewpoints were expressed and their merits hotly debated, economists

have found that income inequality in the United States has steadily increased—implying, perhaps, that we have favored efficiency over equity—and the media have devoted continued coverage to the phenomenon. *The Wall Street Journal* and *The New York Times*, for example, each published major, multipart series earlier this year on inequality trends and their implications.

But the reasons behind growth in U.S. income inequality are poorly understood, and recent research by a Minneapolis Fed economist suggests that the leading theory is an insufficient explanation. Meanwhile, another Fed economist has found that while U.S. income inequality has indeed grown significantly, *consumption* inequality has increased only moderately. The key to this puzzle: deeper use of credit markets. And yet another line of Minneapolis Fed research into economic inequality suggests that Rawls' veil of ignorance as the standard of justice leads logically to a quite surprising outcome: extreme inequality—of result *and* of opportunity.

These three pieces of research have by no means resolved all issues regarding inequality in the United States—any more than the Okun/Rawls/Nozick discussions of the 1970s did. But they *are* shedding light on the right questions to ask and the most promising avenues for finding solid answers.

****

Divergence

There are many ways, of course, to measure the economic well-being of individuals and the differences among them. A June 2003 *Region* article, "Beyond 'Rich' and 'Poor,'" reports on an analysis of such measures, looking at earnings, income and wealth distribution during the 1990s. But in a comprehensive December 2004 *Quarterly Review* article, Minneapolis Fed consultant and University of Minnesota economist Zvi Eckstein, along with Northwestern University economist Éva Nagypál, explores trends in inequality in *earnings*—wages and salary—over the course of four decades.

Analyzing earnings data from the U.S. government's Current Population Survey and contrasting medians (half above, half below) with means (arithmetic averages, skewed by extremes) to quantify inequality, the economists find that mean earnings were 8.5 percent higher than median earnings in 1961. By 1980, the gap had climbed to 18.6 percent, rose slowly over the next 10 years to 19.8 percent and then jumped to 32.6 percent by 2002—a nearly fourfold increase in earnings inequality over 40 years (see graph). Since nonwage benefits such as health insurance and pensions constitute a greater share of labor income for more highly compensated people, Eckstein and Nagypál state that this trend "is probably the lower bound for the trend in compensation inequality."

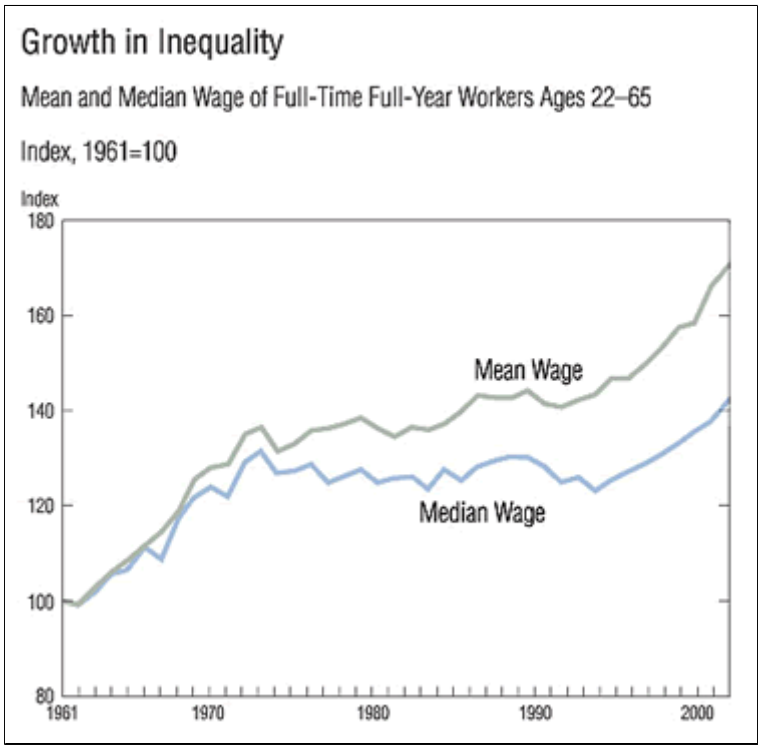

Source: Eckstein and Nagypál, *Federal Reserve Bank of Minneapolis Quarterly Review*, December 2004

The analysis has greater depth when the data are broken down by gender—a point that gains significance as their discussion proceeds. Over the entire period, men's wages increased by 75 percent, but women (partially) closed the gender gap with a 107 percent wage increase. Despite this improvement, significant gender inequality persists: In 2002, women earned only 69 percent as much as men.

Inequality trends also differed for the genders. Among men, inequality in earnings was fairly constant from 1961 to the mid-1970s; it then began a steady increase through 2002. Among women, earnings inequality actually declined from 1961 to 1975 as women shifted out of low-paying jobs, then held steady until 1981. Inequality then started growing consistently through 2002, paralleling the trend among men, though at a lower level.

Why has inequality grown?

The big question, for many, is, What lies behind this steady increase in inequality? And the most popular answer is captured in the phrase “skill-biased technical change.” It's the idea that over the last several decades new workplace technology—computers, in particular—has increased demand for highly educated workers. Their skills are rewarded with higher compensation, whereas workers with less education perform lower-skill jobs and receive much lower wages. As skill-biased technology advances, the wage gap grows.

Despite the difficulty of teasing out the many forces at work in the labor market, the SBTC hypothesis is widely accepted by economists. There's “virtually unanimous agreement,” writes University of Michigan economist George Johnson, in a 1997 *Journal of Economic Perspectives* article, that “relative demand increased for workers at the high end of the skill distribution and thus ... earnings inequality increased.” An influential 1997 Minneapolis Fed research department staff report
(SR 239), later published in *Econometrica*, by Per Krusell et al., buttresses the argument with a quantitative assessment of the skill premium, the wage ratio of high- and low-skilled workers. Their conclusion: “Increased wage inequality results from economic growth driven by new,

efficient technologies embodied in capital equipment."

More recently, the Richmond Fed's 2004 annual report essay, "What's Driving Wage Inequality?" delivers a careful and persuasive case for the theory, referring to "overwhelming evidence and an emerging consensus about the role of skill-biased technical change on the wage structure," while noting continued uncertainty about related phenomena such as declining real wages for less-skilled workers.

In their 2004 analysis, Eckstein and Nagypál provide evidence that supports the SBTC theory. From 1963 to 2002, they find, earnings inequality grew dramatically among different education groups. "This observation is central to the argument that SBTC is the main cause for the observed rise in inequality," note the authors. Among men, the ratio of earnings for postgrads to high school grads nearly doubled between 1963 and 2002 (from 1.4 to 2.6). For women, the inequality growth trend by education group is similar, though less dramatic.

Over the same period, the educational composition of the workforce shifted dramatically toward far higher educational levels. In 1964, 11 percent of female workers had college degrees; by 2003, nearly
one-third were college graduates. The figures are similar for male workers. Taken together, then, the rising wage gap between
college-educated workers and non-college-educated workers and the higher ratio of college- to non-college-educated workers provide underpinnings for the SBTC argument.

Doubts

But other facts raise doubt. Eckstein and Nagypál point out that the lynchpin of the SBTC hypothesis—the education premium—is dominated by the increase in earnings by workers with postgraduate degrees. Yet the percentage of such workers in the overall workforce climbed most dramatically from the mid-1960s to the early 1980s, starting well before "the spectacular rise in their wages," write the authors. The postgrad share of the workforce has been fairly stable since the early 1990s while their wage premium has continued to climb.

Things become more problematic for the SBTC premise when Eckstein and Nagypál analyze inequality trends by occupation. "The SBTC hypothesis posits an increase in the demand for skilled labor," they observe. But when they break down the workforce into three categories—professional, white collar and blue collar—they find that after the early 1980s, there has been very little change in the occupational composition of the male workforce. (From 1983 to 2001, the share of professional workers increased from 30 percent to 33 percent. The share of white and blue collar workers each declined 1.5 percentage points.) For women, on the other hand, the share of the female workforce constituted by professionals increased significantly, but their earnings did not rise dramatically.

Despite this relative stability of occupational distribution, the wage gap between occupations rose over time (at roughly the same rate as the wage gap by education), and the economists found that occupation had significant explanatory power on wage differences—independent of the effect of education. "Therefore," they write, "we believe that any theory addressing the changes in the wage and employment structure should also incorporate occupation as a measure of skill."

Half the sky

The other missing piece, and a major one, is the role of gender. Eckstein and Nagypál point out that the most dramatic shifts in the labor market over the past four decades occurred among women. Their labor force participation climbed dramatically: from 42 percent in 1962 to 72 percent by 1997 and stable since (see graph). (Male labor force participation declined slightly from

94 percent in 1962 to 86 percent in 2003.) The wage gap between men and women declined as women's educational attainment grew. And women experienced less of an increase in earnings inequality over the last 40 years than did men. The SBTC hypothesis, note the economists, has little to say about such changes.

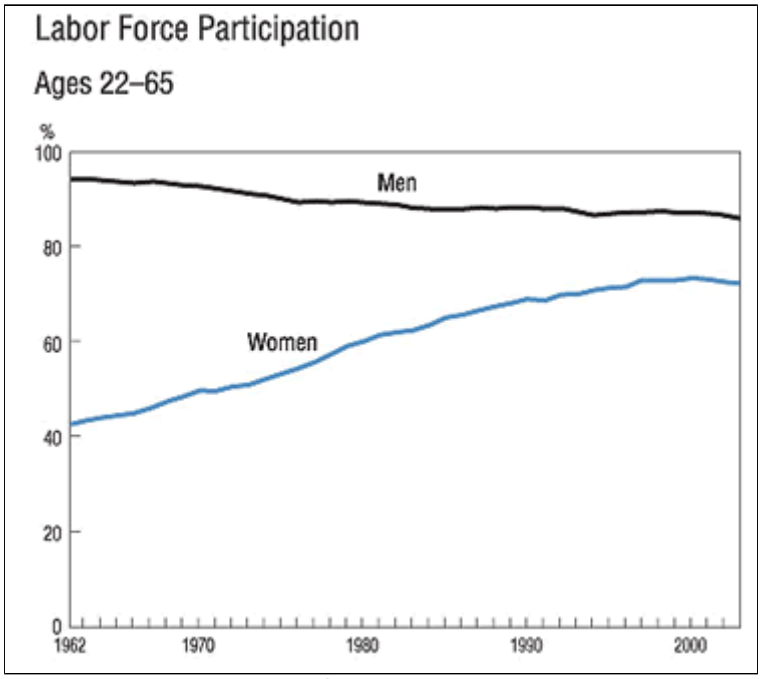

Source: Eckstein and Nagypál, *Federal Reserve Bank of Minneapolis Quarterly Review*, December 2004

Since the SBTC premise relies on the idea that new technology created a demand for highly skilled labor and raised the skill premium for that labor, it provides little room to incorporate or explain the occupational and gender trends that Eckstein and Nagypál uncover, trends that are arguably as significant as the increase in inequality, and quite possibly related to it. “Any theory addressing the changes in the wage and employment structure should also incorporate occupation as a measure of skill,” they write, as well as “an explanation for the dramatic change in women's performance in a labor market.”

The economists, then, are not dismissing the SBTC hypothesis, but rather arguing that it doesn't single-handedly explain the evolution of labor markets and inequality over recent decades. “It is, of course, more likely that there are several factors that account for the main changes that have occurred,” they write. “To understand the importance of the various mechanisms, it is necessary to formulate dynamic models that can quantitatively include the main alternative explanations so that one can measure the impact of each.” (See charts on earnings and occupations for men and women.)

****

You are what you consume

Whatever the explanation(s) may be for labor market evolution and income inequality trends, the real question is, So what? After all, income isn't the ultimate goal for most of us. The payoff is what that money can purchase: food, shelter and Bon Jovi CDs. So the bottom line isn't really how much money we make in a year, but what we're able to buy.

In a recent staff report, “Does Income Inequality Lead to Consumption Inequality?” (SR 363),

Fabrizio Perri, a Minneapolis Fed and New York University economist, and Dirk Krueger of Goethe University Frankfurt, study the empirical and theoretical sides of the question. They begin by confirming the basic income inequality trend.

Examining data from the Consumer Expenditure Survey between 1980 and 2003, they find that, indeed, income inequality has increased significantly. At the beginning of the period, a household whose income fell at the top 10 percent point of the income distribution earned 4.2 times the household sitting at the bottom 10 percent point. By the end, this ratio had surged to over 6. Another common statistical measure of income dispersion—the standard deviation of the logarithm of after-tax labor income plus transfers such as welfare and Social Security income—had climbed 21.4 percent.

And consumption? To their surprise, Krueger and Perri found that the rise in consumption inequality was far less dramatic than the jump in income inequality. The consumption ratio of the richest to poorest grew modestly from 2.9 to 3.4. The variance in log measure rose 5.3 percent. Even with different measures for consumption and different portions of the total population sample, the results remained: Over this time frame at least, a dramatic rise in income inequality had not translated to a similarly large increase in consumption inequality (see graph).

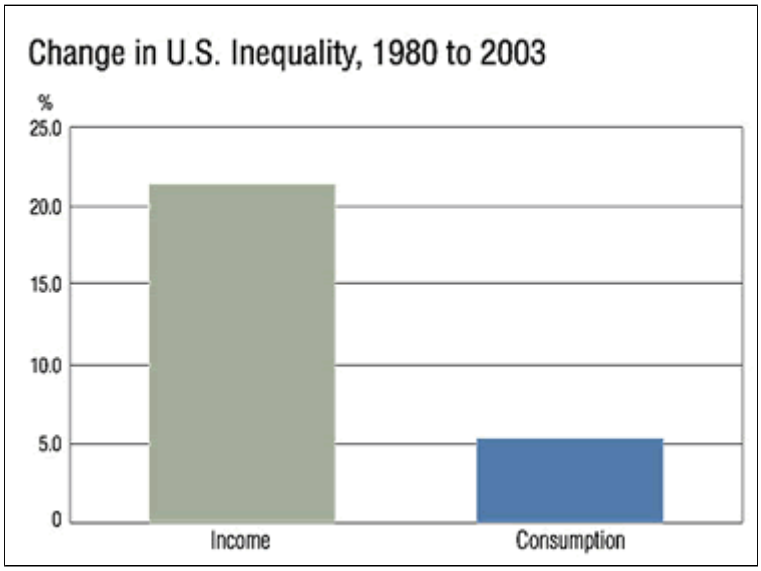

Source: Krueger and Perri, Staff Report 363,
Federal Reserve Bank of Minneapolis, June 2005

When Krueger and Perri analyze the mystery, the major clue appears in remarkably different patterns for “between-group” and “within-group” inequality trends. Essentially, the distinction they're looking at is the general level of income a person expects to attain and maintain in life (due to education, experience, race and gender) versus the year-to-year earnings variation that person might experience due to losing a job or a similar transitory “shock.”

And they find that while income and consumption inequality track quite closely when looking at people's “permanent” income status, the divergence is large when looking at transitory income effects—and the latter account for nearly three-quarters of observed inequality.

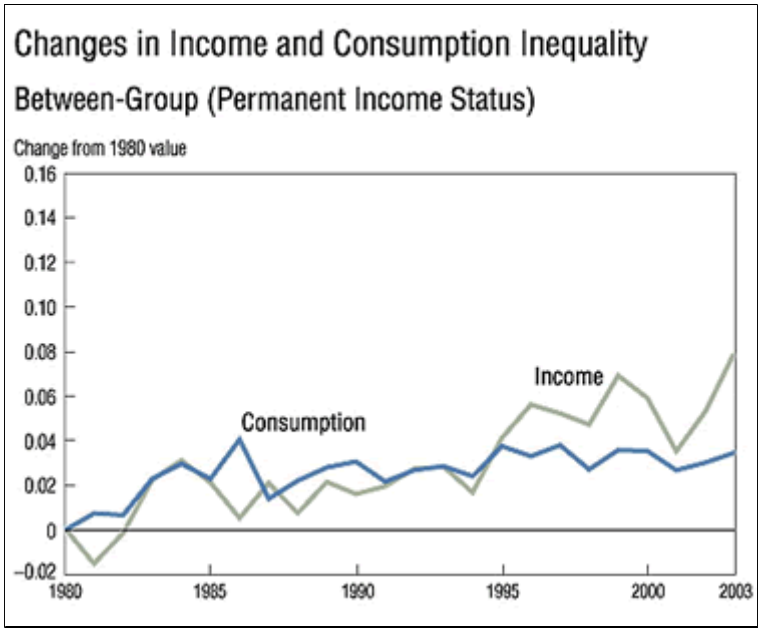

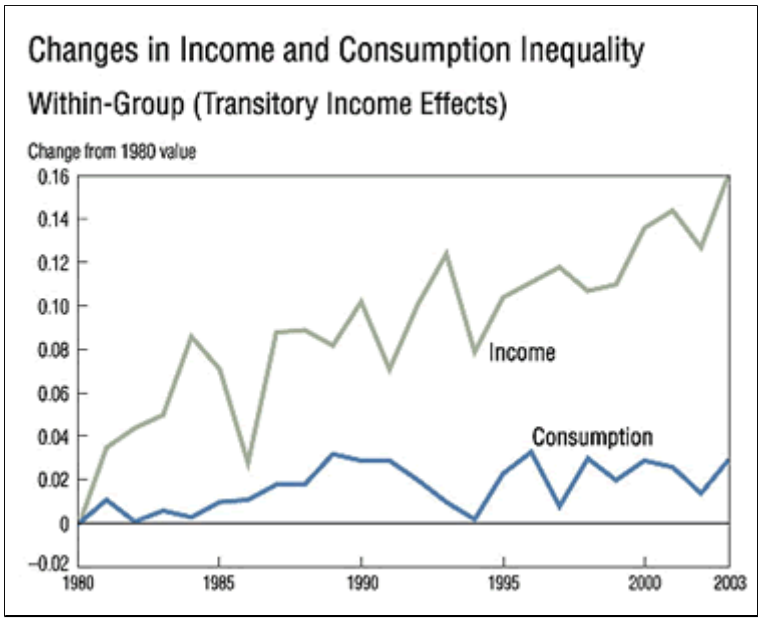

Source: Krueger and Perri, Staff Report 363,
Federal Reserve Bank of Minneapolis, June 2005

On a more intuitive level, Krueger and Perri's idea is that people tend to adjust their consumption patterns to reflect their long-term income prospects; short-term income volatility won't necessarily change spending practices. If a skilled carpenter loses his job but anticipates finding a new one quickly, he won't cancel the purchase agreement on his shiny new pickup.

New shock absorbers

On the other hand, it's not likely that the carpenter will buy that truck without taking out a loan. And that's the key innovation in the model that Krueger and Perri develop to explain why consumption and income inequality trends have been so different. “It is our hypothesis,” they write, “that an increase in the volatility of idiosyncratic labor income ... has not only been an important factor in the increase in income inequality, but has also caused a change in the development of financial markets, allowing individual households to better insure against the (now

bigger) idiosyncratic income fluctuations."

In other words, over the last 20 or so years, household earnings have become less stable, but markets have responded to this volatility by developing credit tools that allow consumers to borrow when they need to and save when they can—smoothing lifestyles with credit cards, home equity loans and the like. If people want to borrow money, whether to pay for college or a large-screen television, financial markets will respond. And on the flip side, people who feel they're doing well (compared to their life expectations) will tend to save, putting something away for a rainy day—"precautionary savings" in the jargon of economics. Financial markets respond there, as well, with a flourishing variety of certificates of deposit and money market accounts.

The economists develop two mathematical models of credit markets to test the hypothesis. The main model assumes that individuals have access to an array of sophisticated borrowing, lending and insurance instruments; in the alternative model, financial markets are much simpler: Households can borrow and lend only at a fixed rate of interest. The main model, they find, slightly understates the actual increase in consumption inequality; the alternative overshoots reality. (While the data show a 2 percent increase in (within-group) consumption inequality from 1980 to 2003, the main model shows a 0.5 percent increase, and the alternative predicts a 4.5 percent rise.)

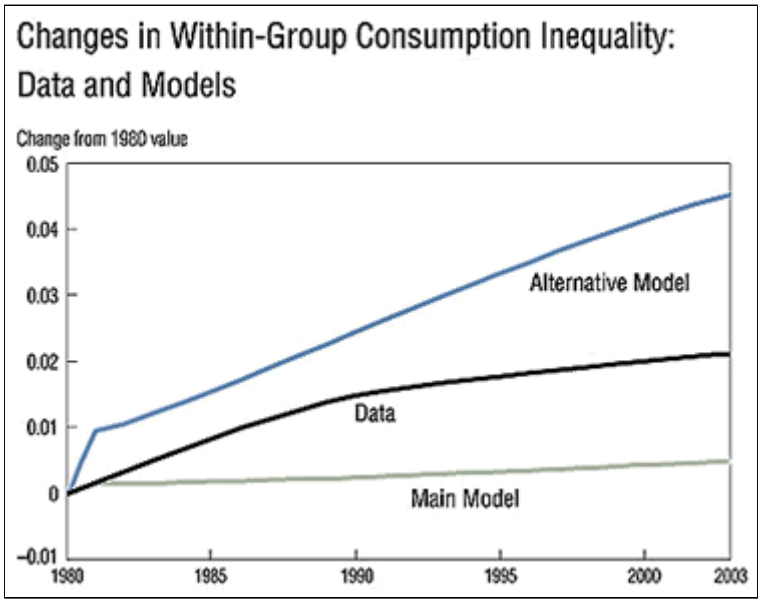

Source: Krueger and Perri, Staff Report 363,
Federal Reserve Bank of Minneapolis, June 2005

Does this suggest that the main model exaggerates the extent to which formal and informal credit markets have evolved to respond to income volatility—that maybe mom and dad won't always come through with the down payment for their kid's first house? Perhaps. The main model "hypothesizes very well developed, maybe too developed direct insurance markets," concede the economists, referring to the ability of individuals to insure themselves against income volatility. In the alternative model, on the other hand, insurance possibilities seem too limited, compared to reality. A future model that assumes something in between, they suggest, "might be even more empirically successful in matching the data."

Regardless, their central point seems solid: Income inequality has surged in recent years, but much of that inequality appears to be transitory—due to temporary volatility—and individuals have, to a considerable degree, smoothed out their peaks and valleys by borrowing and saving to match their anticipated incomes.

****

Back to the veil

If Eckstein and Nagypál have raised doubts about the prevailing wisdom on past inequality, and Kreuger and Perri have suggested that individuals and markets have managed to smooth consumption in order to cope with income volatility, what do we then know about the future of inequality? In three decades, will we see an even greater concentration of earnings, or will the United States experience a fanning out of the income distribution—as general education levels rise, as new technology perhaps favors *un*skilled labor or as women exert still more influence in labor markets?

Predictions of this sort are too speculative to be worthwhile, but theory does tell us something about possibilities. Christopher Phelan, a senior economist at the Minneapolis Fed, has studied the theoretical issues that surround inequality, and his work suggests that high levels of inequality are not only an efficient outcome, from an economist's perspective, but also meet Rawls' criterion of fairness: a choice that would be made from behind the veil of ignorance.

Phelan frames the discussion by looking at how an economy balances the trade-off between insurance and incentives. Societies provide incentives to encourage productive work, but they also offer a measure of security to those who—through bad fortune—are unable to produce enough to survive. The problem, of course: If insurance is too generous, people may decide not to work.

In Phelan's model economy, as in most, individuals try to maximize their utility, related to how much they can consume. The latter is determined by both effort and luck. With hard work and good luck, they'll prosper. But with bad luck, no matter how hard they work, they'll starve. To prevent this, to safeguard its members against misfortune, society provides insurance—a redistribution of the economy's output. And because bad luck can hit anyone, everyone agrees to such an arrangement from behind the veil, even though helping the unfortunate means the lucky, hard workers will have to sacrifice and the economy as a whole may produce a bit less.

What do you know?

This works fine if there's perfect information. If everyone in society is fully aware of how much effort and luck everyone else experiences, they'll be willing to provide generous insurance. “A fair allocation in this economy,” Phelan writes in a 2002 *Quarterly Review* article, “has not only ... equality of opportunity ... but also equality of result—income is redistributed so that all households have the same consumption.”

But what if individuals have “private information”—if people know something about their luck and work effort that society doesn't know? In particular, what if individuals can skip work, pretending to be sick, and still get paid? If people can shirk work without consequence—in economic jargon, the system isn't “incentive compatible”—the economy will find itself with a serious lack of output.

Thus, Phelan's model demonstrates, when private information exists, the only way to get people to work is to provide incentives. And furthermore, when the model runs to its ultimate conclusion, “imposing the incentive constraint implies that ... there is inequality of result.” And the irony is that this result comes about even though all individuals knew it would when they stood behind the veil. “Through no fault of their own, the households with low output realizations have lower consumption levels than those with high output realizations,” writes Phelan. “Why is this fair? It is what all households would have agreed to if they could have chosen in advance.”

Looking forward

The discussion above presents a static scenario—individuals and society start with initial conditions, follow the rules of the social contract and the game ends. But we live—with any luck—in a multiperiod world, and to represent it economists build dynamic models.

"Dynamic contracting theory" creates models in which people play an economic game repeatedly, and the values they start with in period 2 depend on their actions in period 1, *ad infinitum*. Actions have consequences, repeatedly.

In the context of a game regarding incentives, insurance and distribution of an economy's resources, the economic contract is written so that individuals know that their current efforts (or lack thereof) will affect their future resources. If they shirk work and claim insurance in period 1, they'll start with less in period 2. On the other hand, if they work hard now, they'll receive more in the future.

Of course, when you're living in a dynamic world, you've got more than one generation to think about: How will parents' work effort—and luck—affect their children? Phelan shows that in this dynamic model, it's optimal for society to have children's consumption depend on their parents' outcome precisely because it will provide strong incentives for parents to work hard, and their work will significantly increase society's output. "The cost of introducing a small amount of dependency of children's consumption on parents' outcomes is second-order," writes Phelan in September 2005 staff report 323, "while the benefit, or gain, from this dependency (which allows for the better provision of incentives to parents) is first-order."

But the long-run implications are startling.

Economists have worked with such models since the late 1980s and shown that inequality increases *infinitely* under such a dynamic system—not just a few plutocrats and lots of peasants, but one person with all the marbles. As a 1992 paper by Andrew Atkeson and Robert Lucas phrases it: "The efficient allocation delivers an ever-increasing fraction of resources to an ever-diminishing fraction of society's population." And in his 2002 *Quarterly Review* article, Phelan finds that this affects equality of opportunity, as well. "Incentive issues [imply] eventual infinite inequality of both opportunity and result." In other words, not only does the dynamic model suggest that one person will eventually gain all the food, but also all the land used to produce food.

And that's a problem. Not in the moral sense—most economists are loathe to make value judgments—but empirically. Because even though inequality has grown substantially in the last 30 years, and the rich indeed are getting richer, we're still a *long* way from infinite inequality. "It didn't work," said Phelan in an interview. "Once you made these [models] dynamic, you got too much inequality. ... That seemed a conundrum because it didn't seem to describe the world we saw. We don't have exploding inequality. Inequality moves up and it moves down, but it doesn't just continually increase."

Who's behind the veil?

So, what's wrong here? The models show convincingly that individuals standing behind the veil of ignorance will agree to a system of incentives and insurance that results, eventually, in perfect inequality. And they do so because they realize that society will always have more to consume—and they or their children will have a chance of consuming it—if incentives are set up such that consumption in the next period relies on output in the current period. But the model's extreme results don't match reality.

The problem, found Phelan in his 2005 staff report, lies in *whose* preferences are being considered in the social contract. If just the first generation sits behind the veil of ignorance, then perfect inequality is the only possible outcome—even if that first generation is altruistic toward future

generations. The reason, simply put, is that people discount the future—they place slightly less value on something that hasn't happened than on something that has.

So “if you only ask Noah,” says Phelan—referring to the first generation—every subsequent generation will have a slightly greater amount of inequality, and eventual infinite inequality. “But if you actually asked one of Noah's descendents,” he continues—if they stood behind the veil—“*they* would like some insurance. ... They'd be willing to give up some mean consumption in order to lower the variance.”

This, then, is a model in which the social contract agreed to behind the veil provides for no social discounting of the future. It results in insurance, at some level, for bad luck (especially the bad luck of being born to poor parents). But, notes Phelan, it doesn't bring back the full insurance or total equality of opportunity that existed in the static model with perfect information. And significantly, he adds, “You don't get infinite inequality anymore.”

Upstairs, downstairs

Moreover, in this model, when all generations are taken into account directly, you also have social mobility. That is, even though the model results in significant inequality of consumption and opportunity (though short of infinite), subsequent generations will not be locked in permanent “castes.” “The descendents of rich people will eventually become poor,” says Phelan, “and the descendents of poor people will eventually become rich. Dynasties will visit every part of the distribution.”

It's a striking result, a mathematical confirmation of the proverb “from shirt-sleeves to shirt-sleeves in three generations.” And the work has already stimulated subsequent research. After hearing Phelan present an early version of this paper at the Massachusetts Institute of Technology in May 2004, MIT economists Emmanuel Farhi and Iván Werning experimented further. At MIT, Phelan recalls, “I said it's an open question what happens where society discounts but not as much as the people.”

Farhi and Werning explore that question by building a similar dynamic model but with society's discount rate *lowe*r than the private discount rate (rather than *no* social discounting, as in Phelan's model, or equal discounting, as in Atkeson and Lucas'). “This relatively small change produces a drastically different result,” write Farhi and Werning. “[L]ong-run inequality remains bounded, ... social mobility is possible and everyone avoids misery.”

****

An end to metaphors?

Inequality has inspired study by economists, philosophers and other scholars for centuries, demonstrating that it's both a complex phenomenon and a durable one. The work reviewed in this article builds upon previous theory and has inspired further research by other economists. And with every effort we discover something new or confirm an existing truth.

But like the inequality metaphors of the 1970s, each piece of current research examines a crucial aspect of inequality without quite capturing its full dimensions—suggesting that understanding the causes and consequences of economic inequality will remain on the research agenda well into the future.

Top of document