

We Beg to Differ

In an expansive review of a flourishing literature, researchers explore the promise of economic models that allow for human variation

Douglas Clement

Editor

L'homme moyen. The average man. In the mid-1800s, this was a revolutionary concept—the idea that characteristics of individuals in a group could be measured, summed and divided to arrive at a statistical mean, an average, that had real significance and utility. The man who formalized the idea, Belgian statistician Adolphe Quetelet, found average tendencies in a variety of human traits, from chest circumferences to crime rates, and in studies of the human body he developed the Quetelet index, a ratio of weight to height now known as the body mass index, the bane of dieters more than a century later.

Economists in the late 1800s were particularly enamored with the idea of an average man (or woman) since it made descriptions and explanations of economic behavior far simpler and more easily generalized. Alfred Marshall adopted the concept in his 1890 masterpiece, *Principles of Economics*, and gave it a name, “representative,” which implied that an average person or company could stand in for any other in the economy. Marshall’s terms “representative firm” and “representative individual” have been used by economists ever since, generalized to “representative agent.”

For nearly a century, economists used the concept fruitfully. If a macroeconomic model sought to explain relationships among economic variables, it used values aggregated across entire populations composed of representative agents. Average unemployment rates, for instance, were a staple concern of econometricians who, as part of the Cowles

Commission, developed the first large-scale economic models of the United States in the 1940s.

But as Minneapolis Fed senior economist Jonathan Heathcote explains in a recent staff report (available at minneapolisfed.org under “Research”) with co-authors Kjetil Storesletten of the University of Oslo and Giovanni Violante of New York University, the “rational expectations revolution” launched in the 1970s by Robert Lucas, Thomas Sargent and Neil Wallace, among others, “transformed the agenda in macroeconomics.” These economists developed macro models that were grounded in optimal decision-making by *individuals*. In the new models, *individuals* sought to maximize their utility subject to various constraints, and their *individual* actions resulted in aggregate outcomes—inflation, unemployment, interest rates and the like.

Still, the first generation of these quantitative models, shaped by Finn Kydland and Edward Prescott, and elaborated by others over subsequent years, relied on a representative agent, with average values, rather than individuals in all their variety.

Why? “The most important reason for this choice was that economists lacked the tools to solve dynamic models with heterogeneous agents,” write Heathcote, Storesletten and Violante. Developing dynamic economic models that incorporated random variation among a number of economic factors was taxing enough for the formulas and microchips of the 1970s and ’80s. But as computers



L'homme moyen?

Economics fundamentally is about the allocation of resources, and so understanding how different policies affect allocation among individuals is central to economic discourse. Developing models that incorporate those concerns is extremely challenging, but the effort is critical to better understanding key economic issues.

grew more powerful and economists developed more sophisticated mathematical techniques, it gradually became possible to deal with diversity.

“The most important discovery”

Moreover, growing bodies of data collected by labor economists and others made it clear that economic diversity was enormous. “The most important discovery was the evidence on the pervasiveness of heterogeneity and diversity in economic life,” said economist James Heckman in his 2000 Nobel lecture, so it was essential to develop models that could link this evidence to explanatory theory. After all, noted Heckman, “At its heart, economic theory is about individuals and their interactions in markets or other social settings.”

Another Nobel laureate, Kenneth Arrow, made a similar observation in 2004. “One of the things that microeconomics teaches you is that individuals are not alike. ... If we didn’t have heterogeneity, there would be no trade,” he said, and then he added a caveat: “But developing an analytic model with heterogeneous agents is difficult.”

The remarkable variety among people’s econom-

ic situations makes obvious the weakness inherent in relying exclusively on representative agent models—in assuming that everyone is, in effect, average. Economics fundamentally is about the allocation of resources, and so understanding how different policies affect allocation among individuals is central to economic discourse. Developing models that incorporate those concerns is extremely challenging, as Arrow noted, but the effort is critical to better understanding key economic issues.

“There is still a lot you can learn from representative agent models, so for some questions those models are just fine,” Heathcote observed in an interview. “But if you ever have a question where you’re interested in distributional effects, then those models are off the table right away.”

As economists analyze federal proposals that will expand benefits for some and raise taxes for others, for instance, “you need a model where people differ by age, by income or in other ways.” One of the key economic trends of the past several decades is the dramatic widening of the wage structure in the United States. “These trends and their implications for policy and welfare can only be explored within heterogeneous agent models of the macroeconomy,” write Heathcote, Storesletten and Violante.

In their staff report, “Quantitative Macroeconomics with Heterogeneous Households,” the economists review over 180 articles, most of them published in the past decade (indeed, 10 papers are so fresh that they’re cited as “forthcoming”), that have expanded the frontier of economics by introducing diversity at the lowest levels of the economy—the individual household—into models that embody entire national economies.

In their appraisal of this burgeoning literature, the economists give structure to what otherwise seems chaotic growth. They provide a quick background on economists’ initial steps toward heterogeneity, and then order their discussion according to “three themes that are central to understanding how inequality matters for macroeconomics”: efforts to expand understanding of sources of risk and diversity; steps to explore channels of insur-

In Brief

Above average

- Economists have long used “representative agents” to analyze the macroeconomy, essentially assuming that everyone was average. While these models are useful for some purposes, they don’t allow economists to understand key distributional issues.
- In recent years, aided by more powerful computers and mathematics, economists have developed macroeconomic models with “heterogeneous agents”—economic actors that vary.
- In a comprehensive article, three economists review a growing literature on heterogeneous agent models, research that examines different sources of risk, different avenues of insurance and novel outcomes to classic macroeconomic questions such as the cost of business cycles and the cause of the equity premium.

From the moment of conception, every person faces risk and experiences shocks, both negative and positive, that affect their economic status. Economists refer to these as “idiosyncratic” as opposed to “aggregate” risks because each individual has a distinct pattern of life experiences, even though everyone may face the same aggregate oil price shock or national unemployment rate.

PHOTOGRAPH BY MARC NORBERG



Minneapolis Fed senior economist Jonathan Heathcote

ance; and finally, attempts to understand how risk at the individual level interacts with risk aggregated across the macroeconomy.

Accounting for differences

From the moment of conception (being born into a rich or poor family) through childhood and adolescence (receiving a good or bad education) to joining the workforce (getting promotions or losing a job), every person faces risk and experiences shocks, both negative and positive, that affect

their economic status. Economists refer to these as “idiosyncratic” as opposed to “aggregate” risks because each individual has a distinct pattern of life experiences, even though everyone may face the same aggregate oil price shock or national unemployment rate.

To some extent, people can insure themselves against the risks that life presents. They can buy health insurance, life insurance, unemployment insurance. They can start a savings account to tide them over the bad times or buy Treasury bonds (and yes, even stocks) to see them through years of retirement. For the sake of mathematical simplicity, early heterogeneous models assumed that markets were “complete” in the sense that insurance was available against all risk. “Formally, in the literature and the economic models, people understood that there was a lot of idiosyncratic risk,” explained Heathcote. “But if you were willing to assume that you could completely insure that risk, then it was irrelevant for the economy as a whole. You could think of there being a single representative agent and just worry about the aggregate shocks.”

But it was clear, both intuitively and empirically, that complete markets don’t exist. No one, *in utero*, can buy insurance against childhood poverty. Nor can anyone, later in life, insure completely against being fired. Indeed, complete markets would eliminate economic inequality because insurance pools would share all of life’s risks. “If the children of Noah had been able and willing to pool risks ... among themselves and their descendants,” wrote Lucas, “then the vast inequality we see today, within and across societies, would not exist.”

“A more realistic assumption,” said Heathcote, “is that all those insurance markets aren’t there; markets are incomplete. Therefore, people worry a lot about their idiosyncratic risks because you individually are the one who suffers if you get a bad idiosyncratic shock.”

Establishing a standard

Embracing heterogeneity and rejecting complete markets eventually led economists to what

The SIM model includes a large number of individuals who each have distinct levels of productivity and make independent choices of how much to consume, how much to save and even how much to work. Added together, their individual choices determine the total economy's capital, labor supply and prices.

Heathcote, Storesletten and Violante call the “standard incomplete markets” model. The SIM model includes a large number of individuals who each have distinct levels of productivity and make independent choices of how much to consume, how much to save and even how much to work. Added together, their individual choices determine the total economy's capital, labor supply and prices. “The SIM model,” write the economists, “has become a workhorse of quantitative macroeconomics.” It joins together a micro-level perspective of heterogeneous households with a macroeconomic overview that aggregates their activity.

It is, in short, a model of the entire economy. And by plugging in numbers from the vast databases collected on household economic behavior and the broad sets of national accounts, this SIM model enables economists to explore the micro- and macro-level impact of different policy choices: lowering interest rates, for example, or raising taxes on certain households. Most importantly, it allows analysts to understand the *differential* impact of these policy choices on individual households, not just their average effect on the aggregate economy.

As Heathcote, Storesletten and Violante explain, the early generations of SIM models were quite limited. At first, economists using SIM models portrayed households as identical initially; they became heterogeneous only due to external shocks. There was very little room in these early models for free will and individual choice in response to opportunities or policy changes. And households could insure themselves only by buying risk-free bonds. Life insurance policies, multi-earner households and government transfer programs weren't part of the picture.

But economists soon learned that the SIM model was supple—it could be extended to incorporate additional sources of risk, to include more channels of insurance, and to be better integrated into the macroeconomy. These extensions expanded the benchmark model far beyond its original incarnation.

Sources of heterogeneity

Once the door to diversity is opened, the possibilities are endless. In modeling heterogeneous households, economists focused initially on differences in earnings. Even there, however, several options can be considered. Do individuals differ in earnings because of their innate ability, or because of the opportunities (and misadventures) they experience during their lives? No doubt the truth lies in a combination of the two—nature *and* nurture. And given that people have some level of free will, they're likely to respond differently to identical shocks. So over their lifetimes, people will come to differ partly on the basis of different initial conditions, partly because they experience different shocks and partly in their differential response to similar shocks.

Being able to account mathematically for these possibilities is one of the strengths of a SIM model, but it has also led to complex analysis and somewhat indeterminate conclusions. “Separating out how much is predetermined early in life, or is forecastable by agents, versus how much is risk and unforeseen shock, that's tricky. It's a relatively open question,” observed Heathcote.

In looking at earnings, SIM models suggest that it's possible to accurately match the data with models that formulate earnings as a combination of a very persistent component and a transitory component—that is, a strong central earnings tendency that wobbles a bit over time. But sources of the central tendency are subject to debate. One study described by Heathcote, Storesletten and Violante suggests that nearly 90 percent of the differences among people in lifetime earnings is accounted for by factors determined by the time individuals enter the labor market; another suggests that these initial conditions account for less than half the variation.

As Heathcote, Storesletten and Violante write: “Distinguishing between initial conditions and labor market shocks is important, since they have profoundly different policy implications.” To the degree that policies seek to reduce inequality, models that emphasize initial conditions would suggest policies that shape people's potential early in life.

The policy implications of this entire line of inquiry are significant, especially for fiscal policy. Taxes and other government programs may alter individual decisions about saving, about becoming an entrepreneur, getting a college education or searching for a job, and these policies are likely to have measurable consequences for the distribution of income across an economy.

Models that focus on labor market shocks, on the other hand, call for policies that retrain workers who lose jobs because of unlucky shocks. “Examples of both types of policies abound in the U.S. economy,” observe the economists.

Different strokes

While innate characteristics and later shocks have a substantial impact on earnings, economists know that people vary in their reactions to economic conditions. There is a component of the earnings pattern, therefore, that reflects people’s individual decisions about how much to work, how to go about finding a job, how much training and education to accumulate and what occupations to go into. “A substantial portion of earnings dispersion,” write the economists, “may reflect different choices rather than different shocks.”

Several recent papers have explored how people’s income changes as they respond to wage volatility (up or down) by working more or fewer hours. Other research has incorporated “search” models of the labor sector, where workers and employers search for job matches. Inevitably, there is some friction during this process—workers may not take every viable job offer because they hope to find a better one; employers might not hire the first qualified worker. In these models, the standard economic assumption that workers are paid according to their productivity breaks down and the implications can be significant. “Incorporating fully fledged search models of the labor market into equilibrium incomplete markets models is a promising new research avenue,” write the economists.

Another avenue that economists have begun to explore is how workers may respond to wage shocks by opting out of the labor market altogether, though temporarily, by deciding to accumulate additional skills through a return to school or job training. Economists have worked on macroeconomic models that generate inequality in lifetime earnings due to individual differences in initial human capital or learning ability, as well as differ-

ences in unforeseen income shocks. This research also looks at changes over time in the college premium—higher earnings from attaining a college degree—and at how individuals make different decisions to attend college because they face different education costs or different expected returns to a college degree. Still another body of research is looking at how earnings vary because people differ in their tolerance for risk.

The policy implications of this entire line of inquiry are significant, especially for fiscal policy. Taxes and other government programs may alter individual decisions about saving, about becoming an entrepreneur, getting a college education or searching for a job, and these policies are likely to have measurable consequences for the distribution of income across an economy.

Beyond earnings

To understand differential economic well-being, economists have also begun to look at other sources of risk in life, beyond uncertainties in the world of work. Primary among these: health, family and capital. Introducing health shocks into economic models is important for several reasons. Models that incorporate health shocks allow economists to better understand how effectively government programs like Medicare insure households. In addition, health uncertainty affects saving decisions, especially among the elderly.

Another key source of risk in life is one’s family—or more accurately, changes in family situation. Marriage, divorce, the birth of a child and the death of a family member can all have major consequences for economic well-being. In one recent paper, economists include marital status as a risk variable in a SIM model and conclude that it is a greater motivation for precautionary savings than is earnings risk.

Finally, capital. Individuals reviewing their 401(k) statements have seen massive volatility over the past decade in returns to capital, and variation across individuals is also high; some have diversified

Continued on page 40

Losing a job can lead to divorce. A plant closing can cause an individual to lose a job; it can also trigger a fall in local housing prices. To capture all the potential sources of risk faced by individual households, therefore, economists will have to incorporate multivariate systems.

Beg to Differ *from page 35*

and dampened risk levels, but not all. Housing is a primary source of wealth for most households, and house price fluctuations vary widely across regions.

While each of these sources of risk is important in its own right, accurate models of an economy have to account for their interconnections. Losing a job can lead to divorce. A plant closing can cause an individual to lose a job; it can also trigger a fall in local housing prices. To capture all the potential sources of risk faced by individual households, therefore, economists will have to incorporate multivariate systems.

Insurance

True, life is risky. But individuals have developed many forms of insurance to deal with that risk. While the earliest heterogeneous models oversimplified by assuming complete markets—perfect insurance—the next generation oversimplified in the opposite direction, assuming just one form of insurance: risk-free savings bonds. Reality sits between these simplifications—insurance isn't perfect, but it's available from many sources.

One common though less than obvious way to smooth out the financial risks in life is insurance within the family. While most economic models have focused on households with just one earner, recent heterogeneous models have examined how husbands and wives can insure their separate income shocks; indeed, in a 2008 paper, Heathcote, Storesletten and Violante find that as the gender wage gap has diminished, the potential for this form of insurance has risen. "This narrowing gender wage gap has increased the scope for insurance within the family," noted Heathcote. "If one person gets a bad shock and there's another spouse with a similar wage, there's more potential for readjustment of labor supply within the household."

Other economists have looked at intergenerational transfers as a form of insurance. Parents reduce their children's risky future by investing time and money in their health and education—raising future earning potential for the children (and perhaps providing old-age insurance for themselves).

Similarly, young workers may insure themselves against labor market risk by opting to live at home with their parents, who may in turn benefit from an additional source of household income.

Still, families are complex socioeconomic institutions, and it can't be assumed—though many macroeconomic models do—that every member has identical preferences and seeks to maximize overall family utility. A number of recent papers have recognized this and explored models of families with noncooperative intrafamily interaction. "We expect more quantitative work in this area," conclude Heathcote, Storesletten and Violante.

A more obvious source of insurance against income fluctuations is financial markets. Economists have begun to devote attention to models where borrowers can default on debt. "There's a nice set of recent papers looking at bankruptcy and the extent to which that's potentially a useful way to cushion the blow of a bad shock—declaring bankruptcy and starting over," said Heathcote. These models allow economists to explore the implications of alternative bankruptcy laws.

Heathcote, Storesletten and Violante note that easier access to credit in recent years "suggests that households are now better placed to use financial markets to insure against idiosyncratic shocks." Credit cards and home equity loans, for instance, have proliferated over the past decade. "At the same time," note the economists, "financial market innovation may have left the economy more susceptible to aggregate shocks. ... Developing a better understanding of the relationship between access to credit at the individual level and the response of the economy to aggregate shocks is a priority for future work."

Lastly, the economists review recent research on insurance provided by the government in the form of redistributive taxation and social insurance programs. They cite 24 papers that have examined the welfare implications of alternative public policies. To what extent does public insurance crowd out private insurance? What are the distributional impacts of changing the balance of taxation on capital, labor and consumption? Does public education constitute

“There’s a nice set of recent papers looking at bankruptcy and the extent to which that’s potentially a useful way to cushion the blow of a bad shock—declaring bankruptcy and starting over,” said Heathcote. These models allow economists to explore the implications of alternative bankruptcy laws.

social insurance by making redistributive transfers to children from poorer families?

In a 2008 paper, Heathcote, Storesletten and Violante incorporate many of these sources of insurance into a SIM model and show that it can generate results consistent with the path of U.S. consumption inequality over the past 40 years.

From micro to macro

The final section of the economists’ paper examines research on the relationship between individual-level economic risk and risk in the aggregate economy. To some extent “the literature hasn’t done all that much,” admitted Heathcote. “There have been almost these two literatures in parallel. The applied labor, heterogeneous agent micro literature thinking about all these idiosyncratic shocks and what they imply for inequality in different dimensions. And then the traditional macro literature has sort of carried on with representative agent models looking at business cycles and stabilization policy and other traditional macro questions.”

Increasingly, economists are merging these literatures, meshing the micro and macro. And Heathcote, Storesletten and Violante argue that the relationship is quantitatively important: Changing idiosyncratic risk can have a large impact on aggregate economy quantities and prices. Still, said Heathcote, “one big question that’s not really fully addressed is, how do these idiosyncratic shocks and the aggregate ones fit together?”

To some degree, the obstacle has been technical. In standard rational expectations theory, economic agents are said to optimize their utility by forecasting future prices. It’s fairly easy to solve this mathematical optimization problem with representative agent models because individuals are assumed to be identical. But dropping that assumption—incorporating heterogeneity—means solving a numerical problem with an “infinite-dimensional mathematical object.” Suffice to say, infinity doesn’t quite add up.

But 10 years ago, economists Per Krusell and Anthony Smith found a way to cope with this complexity. They showed that in practice individuals in

these models can forecast prices extremely well given a forecasting rule based on a small but carefully chosen set of variables. Their methodological innovation—“approximate aggregation”—allowed economists to explore the relationships between the dynamics of the macroeconomy and the lives of the individuals that constitute it.

Business cycles, inflation and asset prices

This exploration has discovered that some classic macroeconomic problems yield new solutions when viewed through a heterogeneous agent framework.

In 1987, using a representative agent model, Lucas showed that business cycles had relatively little impact on overall socioeconomic welfare. By implication, he proved that eliminating those transitory fluctuations wouldn’t be particularly beneficial. The result seemed counterintuitive in the sense that recessions raise huge public outcry and significant political response.

In recent years, economists have used heterogeneous models to further explore this finding and discovered that under some conditions the Lucas result is overturned. For example, aggregate downturns tend to be times in which idiosyncratic risk—and particularly the risk of unemployment—is particularly large.

Similarly, economists have used heterogeneous models to study the differential impact of inflation. Some studies have found that high expected inflation has a negative impact on the poor because they hold more of their wealth in cash than do the rich. Other research has found that surprise inflation creates large losses for older, wealthy households because they hold more bonds than others.

Another classic problem in macroeconomics: In 1985, Prescott and Rajnish Mehra showed that the large difference in financial returns between stocks and bonds couldn’t be explained under standard economic assumptions. Ever since, economists have been on a quest to solve the “equity premium puzzle.”

Heathcote, Storesletten and Violante review a number of studies that have found partial solutions

"To really get a definitive answer on whether all this heterogeneity matters a lot or not," said Heathcote, "you need a deeper theory of how these idiosyncratic shocks connect to the aggregate world." In other words, economists have yet to develop a cogent explanation of the links between individual economic risk and economywide risk.

to the puzzle with models that allow for heterogeneity. One fruitful approach has been to introduce heterogeneity in preferences or access to financial markets such that a large share of an economy's risky assets ends up in a few hands. Those individuals bearing a disproportionate amount of equity price risk then require a high expected equity premium. Nonetheless, Heathcote, Storesletten and Violante conclude that though incomplete markets and heterogeneity have significant implications for asset pricing, they cannot fully resolve the equity premium puzzle.

More is needed

In conversation, Heathcote noted that while these new perspectives on business cycles and the equity premium are intriguing, they're not entirely satisfying. There are ways to set up a heterogeneous model where business cycles impose large welfare costs, but there are alternative and equally valid setups that don't find large costs from temporary fluctuations.

"To really get a definitive answer on whether all this heterogeneity matters a lot or not," said Heathcote, "you need a deeper theory of how these idiosyncratic shocks connect to the aggregate world." In other words, economists have yet to develop a cogent explanation of the links between individual economic risk and economywide risk. "Economists have noted that there seems to be some correlation, that when things are going bad in the aggregate and the economy is in a recession, it looks as though idiosyncratic risk also gets bigger," he observed. "But precisely why that should be, well, I don't think we have a good understanding of that."

One possible line of explanation has to do with technological change and increased competition that could increase overall productivity and economic growth, but simultaneously raise income volatility and inequality among individuals. For evidence that this explanation has empirical roots, witness the political resistance to outsourcing and lowering trade barriers.

Another promising story that could help economists link individual risk and economywide risk is a

labor or job search theory in which shocks to the economy can lead to waves of job destruction and creation. The idea is that following an aggregate shock that forces some firms to close, there is likely to be a delay in matching workers to new jobs both because the search-and-match process is inherently time-consuming and because workers may need retraining to fit into new jobs.

Whatever the explanation may ultimately be, Heathcote, Storesletten and Violante are convinced that searching for it is a priority, not simply for academic economists but for policymakers. "The main motivation for developing a deeper theory of the interaction between aggregate and idiosyncratic risk," they write, is that "it makes the framework much more useful for policy analysis." Only by moving beyond the abstraction of average people and representative agents can economists "evaluate the distributional impact of aggregate stabilization policies, and the business cycle implications of social insurance policies." **R**

