Wealth Happens

by Mark Buchanan

Harvard Business Review

Reprint R0204C

Harvard Business Review

April 2002

Wealth Happens

Researchers are coming to surprising conclusions about how riches are distributed in societies. Their findings not only have important policy implications but also shed new light on the way complex social and economic networks operate.

by Mark Buchanan

FREE-MARKET THEORIES of economics date back at least to the Scotsman Adam Smith in the latter half of the eighteenth century. In his *Wealth of Nations,* Smith claimed that free trade among the members of a society inevitably leads to an outcome that is good for the society as a whole, even though each individual pursues only his own selfish gain. After all, as he noted, "It is not from the benevolence of the butcher, the brewer, or the baker that we expect our dinner, but from their regard to their own interest. We address ourselves, not to their humanity but to their self-love, and never talk to them of our necessities but of their advantages."

If an individual can profit by manufacturing some product or supplying some service, Smith reasoned, he will do so. And his very ability to turn that profit proves that other members of the society must want those goods or services. In this way, the full spectrum of society's needs will be met through the pursuit of individual self-interest. Such a free-market economy should work smoothly and efficiently without any global management, as if guided and organized by Smith's famous invisible hand.

Today, Smith's metaphor stands at the very center of Western economic thinking. And for more than a century, an army of theoretical economists in the so-called neoclassical tradition has worked diligently to prove that it is indeed true – that individual greed really must translate into collective good. To make their case, they generally begin by assuming that economic agents are not only greedy but also perfectly rational. No one, the theory demands, would be so foolish as to let his emotions get in the way of rational decision making, and no one would simply imitate others without excellent reason. Such assumptions are eminently absurd, of course, and yet there is a method in the absurdity. Once we agree to ignore the vagaries of human behavior, theoreticians

can proceed as if they are dealing with atoms or molecules, and they can prove (or at least try to prove) that the invisible hand really works.

Recently, however, another approach has emerged, one that addresses the complexities of economic reality rather than pushing them into the background to simplify the mathematics. A growing number of economists have embraced the field of behavioral economics, which acknowledges the existence of human irrationality and tries to found economic theories on a more realistic appraisal of people's behavior. Behavioral economists are, for instance, challenging the long-held assumption that stock markets operate in a rational and efficient manner. Economic researchers have demonstrated that stocks which perform extremely well in one year frequently perform poorly in the next, suggesting that investors irrationally tend to overvalue stocks that have gone up in the recent past. Other economists have shown that investors are irrationally predisposed to cling to stocks that have recently lost value, and that the human tendency to follow the pack appears to play an important role in market dynamics. By drawing on such discoveries, researchers have built strikingly accurate models of financial markets in which stock prices fluctuate irregularly, and bubbles and crashes occur as frequently as they do in the real world.

But researchers are also working in another very different direction, stimulated by the realization that in large networks of interacting agents the details sometimes do not matter. To explain the way complicated social and economic networks operate, we may not need to resort to grand ivory-tower theories or complex behavioral models. Broad economic patterns may instead emerge from a few straightforward rules. Whether people are rational or irrational or some delicate mixture of the two, the intricacies of their behavior may have little effect on some of the most basic of all economic realities.

Rich and Poor in America

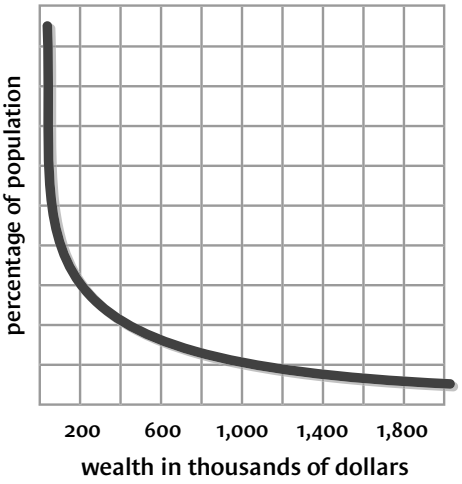

This plot of household wealth in the United States, taken from 1998 census figures, clearly shows a distribution of rich and poor forming a Pareto curve. The highest percentage of households fall at the lower levels of wealth, but at the higher end, the curve drops off relatively slowly, displaying Pareto's "fat-tailed" pattern.

A Universal Law of Wealth

The economic world is full of patterns, many of which exert a profound influence over society and business. One of the most controversial is the distribution of wealth. You might expect the balance between the rich and the poor to vary widely from country to country. Different nations, after all, have different resources and produce different kinds of products. Some rely on agriculture, others on heavy industry, still others on high technology. And their peoples have different backgrounds, skills, and levels of education. But in 1897, an Italian engineer-turned-economist named Vilfredo Pareto discovered a pattern in the distribution of wealth that appears to be every bit as universal as the laws of thermodynamics or chemistry.

Suppose that in the United States or Cuba or Thailand–or any other country for that matter–you count the number of people worth, say, $10,000. Then you count the number of people at many other levels of wealth, both large and small, and you plot the results on a graph. You would find, as Pareto did, many individuals at the lowest end of the scale and fewer and fewer as you progress along the graph toward higher levels of wealth. But when Pareto studied the numbers more closely, he discovered that they dwindled in a very special way toward the wealthy end of the curve: Each time you double the amount of wealth, the number of people falls by a constant factor. The factor varies from country to country, but the pattern remains essentially the same. (See the graph, "Rich and Poor in America" for one example of a Pareto curve.)

Unlike a standard bell curve distribution, in which great deviations from the average are very rare, Pareto's so-called fat-tailed distribution starts very high at the low end, has no bulge in the middle at all, and falls off relatively slowly at the high end, indicating that some number of extremely wealthy people hold the lion's share of a country's riches. In the United States, for example, something like 80% of the wealth is held by only 20% of the people. But this particular 80-20 split is not really the point; in some other country, the precise numbers might be 90-20 or 95-10 or something else. The important point is that the distribution (at the wealthy end, at least) follows a strikingly simple

Mark Buchanan *is the author of* Ubiquity: The Science of History…or Why the World Is Simpler Than We Think *(Crown, 2001) and the forthcoming book about the new science of networks,* Nexus, *from which this article is adapted.* Nexus *will be published by W.W. Norton in North America and by Weidenfeld & Nicolson in the United Kingdom.*

mathematical curve illustrating that a small fraction of people always owns a large fraction of the wealth.

What causes this pattern? Is there some kind of regularity in human behavior or culture that supersedes national variations? Is there some devilish conspiracy among the rich? Not surprisingly, given the strong emotions stirred by matters of wealth and its disparity, economists have flocked to such questions. Of the central issues in economics, John Kenneth Galbraith wrote in his *History of Economics,* the first is "how equitable or inequitable is the income distribution. The explanation and rationalization of the resulting inequality has commanded some of the greatest, or in any case some of the most ingenious, talent in the economics profession." Despite all the attention, however, Pareto's distribution has, from a mathematical standpoint, stubbornly defied explanation.

Finding out why one individual is richer than another is, of course, relatively straightforward. One has only to delve into the details of inheritance and education, inherent ability and desire to make money, circumstance, and plain old luck. The sons or daughters of doctors or bankers frequently become doctors or bankers themselves, while children born into inner-city poverty often remain mired in hardship, unable to escape their environment. But Pareto's distribution isn't about individuals. It captures a pattern that emerges at the level of large groups, leaving individual histories aside. It is, it might be posited, a *network effect.*

As scientists have discovered, many of the overarching organizational features of networks depend only weakly or not at all on the actions or character of their individual members. Physicists, to take just one example, have long known that, in some cases, they can build strikingly accurate models of complex molecular systems using only a few very crude assumptions. It turns out that the details of individual atoms have little influence over the behavior of the entire network. In principle, the same might be true of wealth. Perhaps Pareto's distribution reflects less about people and their characteristics than it does about the deeper, impersonal laws of network organization.

Webs of Wealth

To find out, let's forget for the moment about creativity and risk taking, the distribution of intelligence, and all the other factors that might influence an individual's destiny. Instead, let's focus on the flow of wealth in an economy. Think of an economy as a network of interacting people. At any given time, each person has a certain amount of wealth, and over the days and weeks, that amount will change in one of two fundamental ways. Your employer pays you for your work; you sell your car; you build a patio; you take a vacation in Italy. Such transactions transfer wealth from one person to another along the links in the network. But suppose you purchase a house or a piece of land, and, sadly, its value falls. Or you invest in stocks and, as in the 1990s, the market soars, showering on you a pile of totally new wealth. In such cases, wealth is not merely transferred but actually created or destroyed. Very basically, then, a person's wealth can go up or down either through transactions with others or by earning returns (positive or negative) on investments.

This is hardly news, of course, but it implies that two factors control the basic dynamics in the web of wealth. As people earn salaries, pay rent, buy food, and so on, wealth should flow through the network in a more or less regular way, like water through a network of pipes. Meanwhile, owing to investments, overall wealth should generally increase slowly, even as individuals' wealth randomly kicks up or down as their investments go particularly well or especially poorly.

Obviously, this picture leaves out almost every detail of reality except the most basic. And yet it is intriguing to wonder if these two simple factors might imply something about how wealth ends up being distributed. A couple of years ago, physicists Jean-Philippe Bouchaud and Marc Mézard of the University of Paris took a large step toward answering this question by bringing into the picture one other "obvious" fact – that the value of wealth is relative. A multimillionaire, for example, will not ordinarily sweat losing a few thousand dollars in the stock market, but the same loss would likely be catastrophic for a single parent trying to raise her son while putting herself through college. The value of money depends on how much one already has, and consequently wealthy people tend to invest more than the less wealthy.

With these commonplace observations, Bouchaud and Mézard formulated a set of equations that could follow wealth as it shifts from person to person, as each person receives random gains or losses from his investments, and as those who accumulate more wealth invest relatively more. Equations in hand for a network of 1,000 people, the two physicists set to work with a computer to create an economic model. Not knowing precisely how to link people together into a network of transactions, they tried various alternatives. Unsure of how precisely to set the balance between interpersonal transactions and investment returns, they tried shifting it first one way and then the other. But no matter what they did, the model always produced the same basic shape of wealth distribution – precisely the same shape as Pareto's distribution. This happened even when every person in the model started out with exactly the same amount of money. And it happened when every person was endowed with identical money-making skills.[1]

The finding suggests that the basic inequality in wealth distribution seen in most societies may have little to do with differences in the backgrounds and talents of their citizens. Rather, the disparity appears to be something akin to a law of economic life that emerges naturally as an organizational feature of a network. (For another surprising application of network analysis, see the sidebar, "The Origins of Segregation.")

Shades of Inequality

Bouchaud and Mézard's discovery sug-

gests that the temptation to find complex explanations behind the distribution of wealth may be seriously misguided. What makes wealth fall into the pockets of a few appears to be quite simple. On the one hand, transactions between people tend to spread wealth around. If one person becomes dramatically wealthy, she may start a business, build a house, and consume more products, and in each case wealth will tend to flow out to others in the network. Conversely, if a person becomes terribly poor, he will tend to purchase fewer products, and less wealth will flow through links going away from him. Overall, the flow of funds along links in the network should act to wash away wealth disparities.

But it seems that this washing out effect never manages to gain the upper hand, for the random returns on investment drive a counterbalancing rich-get-richer phenomenon. Even if everyone starts out equal, differences in investment luck will cause some people to start to accumulate more wealth than others. Those who are lucky will tend to invest more and so have a chance to make greater gains still. Hence, a string of positive returns builds a person's wealth not merely by addition but by multiplication, as each subsequent gain grows ever bigger. This is enough, even in a world of equals where returns on investment are entirely random, to stir up huge wealth disparities in the population.

That doesn't mean that inequities in wealth can't be mitigated. In a Pareto distribution, the factor by which the number of people declines as wealth increases remains constant in any particular country, but the factor itself is different in different countries. So, while there is always a disparity between the rich and the poor, there are differences in degree from country to country. And, socially speaking, there's a world of difference between an 80-20 distribution and 90-5.

Bouchaud and Mézard's network model can track those degrees of inequality and show how Pareto's distribution can be influenced. Specifically, the two researchers found that the greater the volume of money flowing through the economy and the more often it changes hands, the greater the equality. Conversely, the more volatile investment returns are, the richer the rich tend to get.

The model was never meant to provide detailed recommendations for public policy. But it nonetheless has some clear implications. Take taxes, for instance. The model confirms the assumption that income taxes will tend to erode differences in wealth, as long as those taxes are redistributed to the society in a more or less equal way. After all, taxation represents the artificial addition of some extra transactional links into the network, along which wealth can flow from the rich toward the poor. Similarly, a rise in capital gains taxes will also tend to ameliorate wealth disparities, both by discouraging speculation and by decreasing the returns from it. On the other hand – and this is somewhat more surprising – the model suggests that sales taxes, even those targeted at luxury goods, might well exaggerate differences in wealth by dampening purchasing (thus reducing the number of transactional links) and encouraging people to invest more of their money.

The model also offers an excellent test of some arguments that politicians use to justify various policies. In the United States, for example, the 1980s and 1990s were dominated by free-market ideology and government deregulation, much of it defended by the idea that wealth would trickle down to the poor. Everything was done to encourage investment activity of all kinds, regardless of the risks involved. It was no accident that this was the era of junk bonds, the savings and loan debacle, and the dot-com boom. Did the wealth trickle down? Under the network model, one would expect just the opposite. A dramatic increase in investment activity, unmatched by measures to boost the flow of funds between people, ought to lead to an increase in wealth inequality. And, indeed, that's exactly what happened. Today, wealth distribution in the United States is significantly less equitable than it was three decades ago. Wealth is more highly concentrated than it is in European nations, verging on the level seen in Latin American countries.

Wealth's Tipping Point

In further studying their model, Bouchaud and Mézard made another, truly alarming discovery. They found that if investment returns grow sufficiently volatile, they can completely overwhelm the natural diffusion of wealth generated by transactions. In such a case, an economy can suddenly reach a tipping point, and wealth, instead of being held by a small minority, will condense into the pockets of a mere handful of super-rich robber barons. The fat-tailed curve will collapse into a very thin-tailed curve even though Pareto's formula will still hold true.

This may not sound dramatic, but in a society of millions it would be. The wealthiest 10% of the U.S. population is a group of 300,000 people, and the collapse of their wealth into the hands of just five or six individuals would represent a dramatic transformation of society. An attendant power shift could follow, with potentially great political ramifications.

As worrying as this scenario might be, it is not science fiction. Though the network model is abstract, its abstraction is also its advantage, for it indicates mathematically, with few disputable assumptions, that a tipping point of this sort must exist in any economy. The U.S. economy may currently be far from this point or close to it. No one knows. In any case, policy makers ought at least to be aware of the precipice over which an economy might tumble.

It is intriguing to wonder if some countries, particularly developing nations, may already be on the far side of the tipping point. It has been estimated, for example, that the richest 40 people in Mexico have nearly 30% of the wealth. It could be, also, that many societies have gone through this phase in the past. Long-term economic trends during the twentieth century lend some credence to this idea, as the total share of the richest individuals in England, for

example, has fallen during the last century, particularly between 1950 and 1980.

Political instability may cause an economy to plunge into this phase. In Russia, following the collapse of the Soviet Union, wealth has become spectacularly concentrated; inequality is dramatically higher than in any country in the West. No one can be sure why, but the model would suggest that both increased investment volatility and lack of opportunities for wealth redistribution might be at work. In the social vacuum created by the end of the Soviet era, few regulations protect the environment or safeguard workers, and so economic activity is less restricted than in the West. This not only leads to pollution and human exploitation but also generates extraordinary windfall profits for some companies. Economists have also pointed out that Russia has been slow to implement income taxes that would help to redistribute wealth.

Again, this simple model is not the final word in explaining the distribution of wealth or how best to manage it. Nevertheless, by starting with remarkably simple assumptions and then studying the patterns that emerge in a network of interacting agents, Bouchaud and Mézard have succeeded in explaining one of the most basic patterns ever observed in economic life.

And this is just one example of the way network analysis may reshape our understanding of economies. At the individual level, to be sure, few details of economic life may be predictable. It may be just as impossible to tell which person will end up wealthy or which business will succeed as to foretell the movements of individual stocks on any given day. But the broader patterns of economic law arise at the level of many people or many companies or in the statistics of price fluctuations over the longer run. As it turns out, both company size and price fluctuations follow patterns very similar to Pareto's distribution. If you look at the total change in the price of a stock over a year, for instance, you'll find that a large fraction of that change will have occurred in just a small fraction of the elapsed time. Many other economic phenomena follow similar patterns, and they can be explained as network effects with similar simplicity. The workings of even the most vast and intricate systems may not be anywhere near as complicated as they first appear. ▽

The Origins of Segregation

The study of network effects doesn't just illuminate the workings of economies. It can also shed light on social networks and their problems. Consider, for example, racial segregation. In the United States, the persistence of segregation is usually attributed to racism or to biased government policies or real estate practices. But another, less obvious factor may be equally influential, as Harvard political economist Thomas Schelling pointed out in the *Journal of Mathematical Sociology* more than 30 years ago.

In trying to uncover the origins of segregation, Schelling began by imagining a society in which most people truly wish to live in balanced and racially integrated communities. There was just one stipulation: Most people would prefer not to end up in a small minority. A white man, for instance, might have black friends and colleagues, and he might be happy to live in a predominantly black neighborhood. But he might not want to be one of the *only* whites living there. This is hardly a racist attitude and may indeed be one that many people share – blacks, whites, Hispanics, Asians, or what have you.

But innocent individual preferences of this sort can have startling effects on a large network, as Schelling discovered when he drew a grid of squares on a piece of paper and played an illuminating game. On his grid, he first placed at random an equal number of black and white pieces depicting an integrated society of two races mingling uniformly. He then supposed that every piece would prefer not to live in a minority of less than, say, 30%. So, taking one piece at a time, Schelling checked to see if less than 30% of its neighbors were of the same color, and, if this was the case, he let that piece migrate to the nearest open square. He then repeated this procedure over and over until finally no piece lived in a local minority of less than 30%. To his surprise, Schelling discovered that the black and white pieces had not just become less uniformly mixed but had in fact come to live in entirely distinct enclaves. The desire to avoid being a member of a small minority had the paradoxical but inexorable effect of obliterating mixed communities altogether.

Schelling's exercise doesn't show that racism or unfair institutions have nothing to do with perpetuating segregation. But it does reveal how surprising kinds of organization can well up naturally in complex networks and how imperative it is to look beyond the individual components of a network in trying to comprehend such effects. Even if every trace of racism were to vanish tomorrow, there may still be a natural tendency for races to separate. Social realities are fashioned not only by people's desires but by more or less mechanical forces – in this case, forces that can amplify slight and seemingly harmless personal preferences into dramatic and troubling consequences.

1. Jean-Philippe Bouchaud and Marc Mézard, "Wealth Condensation in a Simple Model of Economy," *Physica,* February 2000. Other researchers have explored similar models and reached very similar conclusions. See, for example, Ofer Malcai, Ofer Biham, and Sorin Solomon, "Power-Law Distributions and Lévy-Stable Intermittent Fluctuations in Stochastic Systems of Many Autocatalytic Elements," *Physical Review E,* August 1999.

Reprint R0204C

To place an order, call 1-800-988-0886.