The information revolution has had a profound impact on the economy but very little impact on economic policy, which is largely still generated by 19th-century ideas.... Having a 21st-century economy based on laws designed for the manufacturing sector is really quite ridiculous.

Economic Policy in the Information Age

by Simon J. Wilkie

The standard economic model, in which many buyers and many sellers compete with one another in the marketplace, works well for tangible goods such as bananas. But you get into deep trouble when you try to apply it to network commodities such as electricity.

The United States is in the throes of the third industrial revolution—the first one, of course, being the harnessing of mechanical power, and the second one being the harnessing of electrical power. This one is the harnessing of information, and has been sparked by the biggest capital investment in the history of humankind. The U.S. spent roughly $4 trillion on information technology, broadly defined, from 1960 to 1994, and it's expected that we'll be spending a trillion a year by 2005, even with the current slowdown. Like the other industrial revolutions, it has taken 30 to 50 years for the results to show up. The productivity gains we've seen in the economy in the last few years have their origins in the early investments that are just now starting to kick in.

The information revolution has had a profound impact on the economy but very little impact on economic policy, which is still largely generated by 19th-century ideas. The legal framework was set by the Sherman Act and the Clayton Act, which were developed in the 1890s to bust trusts such as Standard Oil. The notion of regulating utilities appeared in the early part of the 20th century, culminating in the Telephone Regulation Act of 1933, which established AT&T as a monopoly. Manufacturing, however, which was the dominant paradigm at the turn of the century, is now less than 17 percent of our economy. It's going the way of agriculture. Health care is now almost 15 percent of the economy, and in a couple of years, it's going to be bigger than manufacturing. Having a 21st-century economy based on laws designed for the manufacturing sector is really quite ridiculous.

There are two main points I want to make. The first is that minutiae are important in the design of economic institutions—that is, the details matter. And they matter a lot. Which is kind of ironic because most economic-policy debate is big-think debate: should we have a market or not? Should an industry be regulated or unregulated? These are the wrong questions. The important questions are really in the small details of how a market is structured. Second, I want to make a case for the fundamental importance of abstract economic theory—the type of arcane research, divorced from the real world, that we do here at Caltech. If you don't pay attention to these extremely mathematical, abstract models, you're bound to make disastrous policy mistakes. So minutiae are important, and the boring questions are really the interesting ones. I'm going to apply this perspective to the Microsoft antitrust case and to California's electricity-deregulation debacle, which is an endless source of fun until your bill comes at the end of the month.

To see where we need to go, we first need to know where we are. The standard economic model says we have a market in which many buyers and many sellers compete with one another. Each buyer has a personal valuation for each good. When you buy bananas, you know how much you're willing to pay for any given quantity. And each seller knows his or her production costs, and sets an individual minimum price accordingly. The market adjusts supply to be equal to demand, as that old axiom says. This is summed up on the next page, in the economist's favorite diagram. The vertical axis is price, and the horizontal axis is quantity. The downward-sloping demand curve says that the lower the price is, the more of a product people are willing to buy. The supply curve slopes upward, saying that the higher the market price, the larger the quantity that sellers will want to sell. At the intersection lies the equilibrium, which is the quantity (q^*) that is actually traded at the market price (p^*).

This competitive market has several desirable properties, the first of which is efficiency, or **E**. That is, the market maximizes the sum of the valuations minus the sum of the costs. This maximizes social welfare—the greatest good for the greatest number. (However, it doesn't mean

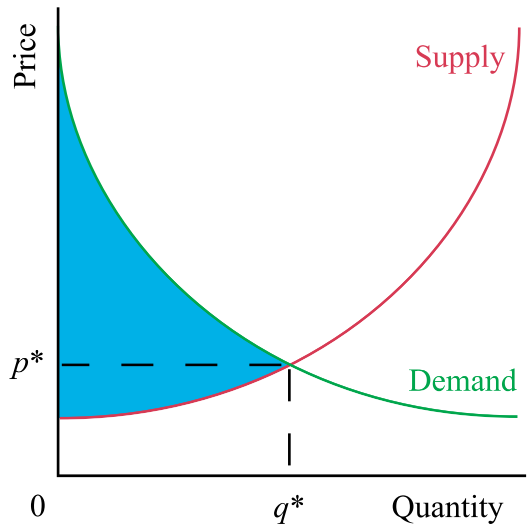

The competitive-market supply and demand curves as drummed into the heads of generations of business majors. The supply curve (more properly, the marginal cost curve) rises while the demand (or marginal value) curve falls. They intersect at point (p^*, q^*), and the shaded region is the social surplus, or value in excess of cost, that people get from trading in the market.

it maximizes *your* welfare. For that, you'd need to do a weighted sum, where your weight was higher than anybody else's.) The distance between the marginal cost (the cost of producing one more unit of the product) and the marginal valuation (the price the consumer is willing to pay for one more unit) is the amount of what we call surplus, or benefit, that people get from trading in the market. Mathematically, the market starts at quantity $q = 0$, and runs until q equals q^*, at which point the market price has dropped to p^* and the sellers are selling at cost, so they quit. But they made a profit on all the previous sales.

The telephone system, the Internet, and the power grid are obviously networks. But HMOs are actually networks, too.... A country club is a network. When you join the club, you get the use of their facilities; you also get to enjoy, or unenjoy, the company of the other members of the club. In fact, another word for "schmoozing" *is* "networking."

And at p^* the buyers, who have been buying the product for less than their personal valuations, are paying as much for the last unit of the product as they think it's worth, so they quit. But they got a bargain on all the previous units. So the market is efficient because it maximizes the integral—the shaded region between the two curves—without even knowing what that integral is. Adam Smith discovered the magic of this "invisible hand" a couple of centuries ago.

The market has three other remarkable properties that we don't talk about as much. First, notice that there's no Tony Soprano—we don't need coercion to get people to use the market. Participation is voluntary, because until the market hits equilibrium and shuts down, every buyer leaves with a bargain and every seller leaves with a profit. We call this property **V**. The next property is **B**, for balance. Supply equals demand without needing an infusion of cash or goods external to the system—unlike, say, Russia, which is kept afloat by large amounts of funds flowing in from the World Bank and the United States. Or your kids might trade toys, but that only works so long as there's a perpetual infusion of new toys from the parents. But a normal adult market has the miracle of always being balanced. The last property, **S**, is the most important one—this market is strategy-proof. People have no incentive to game the market. If I go to Von's supermarket to buy three pounds of bananas, and they're a dollar a pound, I have no incentive to buy four pounds, and I don't think, "Ha! I'll fool them and only buy two pounds!" If I want three pounds, I buy three pounds. There is absolutely no benefit to me from strategic behavior. To sum up, the competitive-market model has four really nice properties: it's efficient, it's voluntary, it's balanced, and it's strategy-proof. That's why policymakers tend to be opposed to monopolies and market regulation, which short-circuit the market's functioning.

But the general paradigm for the 21st century is a network-market model. Network markets are ubiquitous—the telephone system, the Internet, and the power grid are obviously networks. But HMOs are actually networks, too. When you join a primary-care physician's group, you're also signing up for the set of specialists affiliated with that group. Banks and ATMs are networks. A country club is a network. When you join the club, you get the use of their facilities; you also get to enjoy, or unenjoy, the company of the other members of the club. In fact, another word for "schmoozing" *is* "networking." The eBay Web site is a network. Network markets work precisely

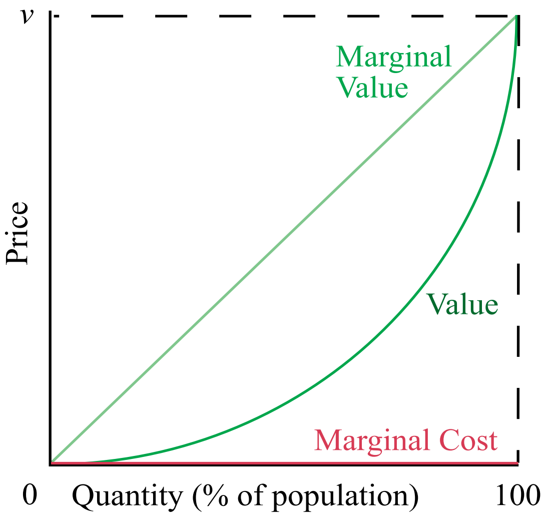

Things look quite different in a network market. Now the marginal-cost curve is flat, at next to nothing, while the product's value increases as more people buy it. The marginal value, which is the derivative of the value, increases even faster until 100 percent of the population owns the product and its value is maximized. This means that a monopolistic seller can name any price up to the maximum value, *v*, and people will still buy the product. In this case, a monopoly can be efficient, in the economic sense of the word.

because of the mass of users that they attract.

In a network market, neither the supply curve nor the demand curve behaves as expected. A supply curve only slopes upward when the marginal cost, the derivative of the cost function, is increasing. And the marginal cost *does* go up, for traditional commodities. Most people think of the marginal cost as going down as more units are produced, but that's only true up to a point. If I'm growing bananas, I'll cultivate my most fertile land first. As demand grows, I'll use increasingly poorer land, and I'll have to buy more fertilizer and more water to produce a crop. The same is true of steel—if a blast furnace runs around the clock, labor costs will skyrocket. The foundry either has to hire more people or pay massive overtime. So there's actually a dis-economy of scale. But let's think about software for a moment: The marginal cost to Microsoft of me buying an extra copy of their operating system is pretty much zero. They have the fixed cost of developing the product, and then the cost of burning one more CD is virtually nil. That's also true of the Internet—the cost of setting it up was huge, and the cost of adding an extra unit is minuscule by comparison. Most network-structured economies have this fundamental problem that supply tends not to be upward-sloping. They really *do* make it up on volume.

Even more troublesome, the demand curve slopes downward only when the *quality* of the product is inherent and is independent of its *quantity*. But network commodity's quality is systemic; it's not inherent in the commodity itself, as it is in a banana. In the electricity market, if I decide to flip on my air conditioner, it affects the voltage—minusculely, but it affects the quality of the service that everybody else in my neighborhood gets. And a network's value to a user also depends on its quantity, but not on the quantity you buy as an individual—rather, the value depends on the number, or the identity, of users. If I'm the only person with a telephone, it's worthless. Its value to me increases as other people buy it, because then I have more people I can talk to. And if I were the only person in the country club, I wouldn't be willing to pay very much to join it. The value typically increases with the number

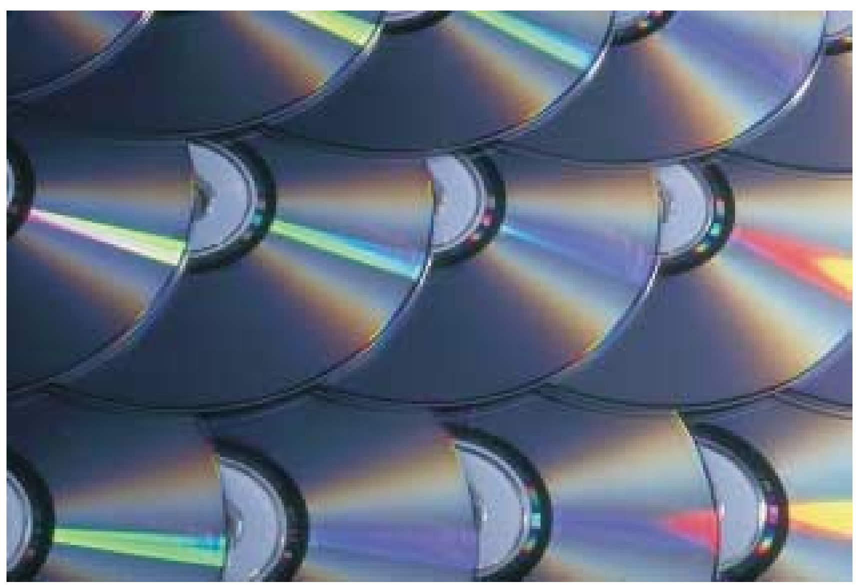

Von Neumann thought he was the smartest person in the world, so he couldn't understand why he kept losing at poker. Being von Neumann, he decided to figure it out, and he developed the theory of strategic interaction between individuals.

of users, but not necessarily. If you have a cable modem, your downloading speed is divided by the number of users who share your connection. Now you have a positive benefit proportional to the total number of users, plus a negative benefit proportional to the number of users who live on your block. The net benefit, if you'll pardon the pun, of having your neighbor in the system may be negative. We call such attributes externalities, because now the product's value is external to the product itself.

So we have to throw our beloved picture out and go to a more abstract framework to analyze network markets. We use game theory, a mathematical technique developed in the 1940s by John von Neumann and Oskar Morgenstern at Princeton. Von Neumann thought he was the smartest person in the world, so he couldn't understand why he kept losing at poker. Being von Neumann, he decided to figure it out, and he developed the theory of strategic interaction between individuals. Game theory lay dormant until the 1970s, when it was seized upon by economists starting to get interested in the economics of information. In fact, a lot of the pioneering work on game theory in economics was done here at Caltech by some of my colleagues—Matt Jackson, John Ledyard, Dick McKelvey, Tom Palfrey (PhD '81), and Charlie Plott. When we model the market as a game, we ask: can we design an economic mechanism, like we would design an engineering device, that has the attributes we want and solves the problem we want to solve? Is it mathematically possible to construct such a thing?

The network math works like this. We have a set of alternatives, which we call **A**'s. For the electricity grid, each **A** would be a possible topology of the network: the capacities of the transmission lines, how they are connected, and so on. The benefit I get from being in this network depends on the choice of **A**—if we're talking about the Internet, I'd like a high-speed connection better than a low-speed connection, for instance—and it depends on **U**, the particular group of users. It might depend positively on **U**, as in a telephone network; or it might depend negatively on **U**, as in the cable modem example, or in a swimming pool—the more people in the water, the larger the negative impact when I jump in.

We can describe our four goals mathematically. Efficiency requires us to choose the **A** and the **U** that maximizes the sum of our individual welfares, minus the cost of providing network configuration **A**. Voluntariness means that if we pay for the network by charging each user an amount, T_i, which could be a flat fee or a function of some sort, the benefit you get minus the money you pay has to be nonnegative. In other words, everybody comes out ahead from being part of the network, or at least breaks even. And the network itself breaks even—that's balance. The sum of the T_i's has to equal the cost of the network. If the T_i's are insufficient, we need an external infusion of cash or assets. If they exceed the cost, we have to decide what to do with the surplus. And finally, strategy-proofness says that your benefit from (**A**,**U**) minus T_i has to be at least as great as the benefit you could get by lying, by manipulating the system to induce some other (**A**′,**U**′) minus your T_i' for that different choice of network conditions and users. For instance, if I'm in a rural area, it might be very expensive to connect me to the network. But if I lie and say I have an enormous value, then it's still efficient to connect me, so it's in the interests of rural users to manipulate their values upwards to ensure that they're connected. The only way to stop that from happening is to make their T_i's extremely high. In fact, U.S. policy is exactly the opposite—we subsidize rural users to help them connect to the network.

Game theory has led to several relevant theorems. The fundamental one, published independently in the early 1970s by Allan Gibbard (then

It's a holdup in both senses—it's an impediment to your use of the network, and it's highway robbery.

at Chicago, now at Michigan) and Mark Satterthwaite (BS '67) at Northwestern, stated that it's impossible to find a mechanism that satisfies our four requirements that it be efficient, voluntary, strategy-proof, and balanced. However, the approach was so abstract that the possibility remained that for some class of network models one could, in fact, have all four. This hope was dashed in 1979, when Harvard's Jerry Green and J. J. Laffont (Laffont is now at the University of Toulouse) revisited the issue. Their work was done in the context of providing a public good, such as building a bridge, but applies to networks as well. It says that no general network-market model can satisfy all four of our desiderata.

Caltech's Matt Jackson, then at Northwestern, in collaboration with Salvador Barberà at the Universitat Autonoma de Barcelona, found that if we're willing to chuck out our beloved, slavish devotion to efficiency, we can come up with a mechanism that will satisfy the other three requirements. We won't need the Sopranos, people won't game the system, and it requires no external infusion of funds. However, the mechanism looks a lot like price caps, which makes industry very nervous. This result was also obtained independently by Hervé Moulin (then at Duke, now at Rice) and Scott Shenker (a computer scientist then at Xerox PARC, now at Berkeley, who was interested in network protocols). And a theorem by Ted Groves at UC San Diego says, when applied to this context, that we can keep efficiency while getting strategy-proofness and voluntary participation, if we're willing give up balance.

And finally, several other people and I have shown that we can get efficiency, voluntary participation, and balance if we're willing to give up strategy-proofness as a global concept and replace it with a local concept. That is, instead of it not being in anybody's interest to game the market *ever*, it's not in my interest to game the market as long as nobody else is gaming the market. If everyone else is playing fair, the system enforces fair play on my part. But if a group of people collude and try to game the market, they can do it. This local strategy-proofness is called the Nash equilibrium, because mathematician John Nash developed the idea at Princeton in about 1950.

So the state of the art in network models is that we can get three out of four. It's mathematically impossible to achieve all four. This means that we have to tailor each market to the particular characteristics of the network it serves—there is no one-size-fits-all optimal policy. And yet, our economic policy is still largely driven by the standard, competitive-market goal of four out of four. But in fact, when we look at the best way to handle different network markets, we may arrive at conclusions that are polar opposites of each other, as I'll demonstrate with a couple of examples—Microsoft and electricity.

Microsoft's operating system is a network externality, because the more people that use it, the more products are developed for it, and the more benefit you get from it. (I use a Mac, myself, so I'm denied a bunch of software that other people have; but for some reason I get more benefit from having a Mac, and fewer friends, than other people do from enjoying more friends and cheaper products. Go figure.)

The problem with this network is a really subtle one, but it's very interesting. It's what economists call the holdup problem, and it occurs when you have a network made of different components that are priced separately. Imagine that you're in New York City and you want to travel down to Washington, D.C., to protest some issue. You jump in your car and you get on the New Jersey Turnpike, which is a toll road. You drive through Jersey to the Delaware Pike, you pay a second fee to Delaware, and you get to Washington. Here's the holdup problem: suppose it's worth a dollar to you to take the trip, and the Jersey Turnpike charges 50 cents. A dollar minus 50 cents leaves 50 cents. You drive down the road and you get to the Delaware Turnpike. What if the Delaware Turnpike hits you for a buck? They have you over a barrel—you've already spent 50 cents, but you're closer now than ever and it's still worth a dollar to you to finish the trip. So you go on to Washington, you eat the other 50 cents, and you mutter to yourself, "What a rip-off. I'll never do that again." So it's a holdup in both senses—it's an impediment to your use of the network, and it's highway robbery.

Now, imagine I'm Microsoft and I've got a nice little monopoly going, with all these peripheral products adding value to my network, and then somebody new comes along with something as essential as my operating system. Suddenly, in order to get the full benefit of your computer, you have to buy my product plus this other guy's product. But you have to buy *my* product first.

A browser isn't much good without an operating system to run it on—at the moment. This means that the industry overall faces a holdup problem. Microsoft knows that if it charges a high price for its product, then the browser company, which now has a captive market, can also charge a high price. But then nobody will buy either product. So the first firm in—I hate to use this expression—the value chain really *is* threatened by the firms farther down the line. In the big picture, it might actually be mathematically efficient for the first firm in

Let's move on to the mother of all mess-ups: electricity "deregulation." I've got that in quotes because people usually think of deregulation as removing regulations, but this "deregulation" produced a new set of rules the size of a phone book.

line to kill off the second firm and integrate the two products. Predatory pricing—selling the second good for free, or below cost, in order to kill a competitor—is illegal, but it solves the holdup problem. Maybe Microsoft *should* be allowed to decide whom to subsidize and whom to kill, given that it already has the operating-system market tied up. If Microsoft had a viable competitor in that market, that might not be true, because then people would have an alternative route to make the journey.

Current federal policy is driven by the idea that we want to break Microsoft's monopoly because monopolies are bad. We outlaw predatory pricing, because predatory pricing enhances monopoly. We want to have open systems, open network platforms, to guarantee the largest amount of access and the largest amount of product development. But in the network model, none of those things can be shown to always be efficient. Microsoft's strategy has encouraged innovation in some areas and thwarted it in others, so the overall effect on efficiency isn't clear. We don't know what the efficient policy actually is, and it might not be that the efficient policy is the best policy. For example, if the efficient mechanism's not balanced, the social cost—the flow of money in or out of the system—might outweigh the benefits. It's a very complex issue, and we need to spend a lot of time modeling the minutiae of the industry in order to get the right solution.

Let's move on to the mother of all mess-ups: electricity "deregulation." I've got that in quotes because people usually think of deregulation as removing regulations, but this "deregulation" produced a new set of rules the size of a phone book. Electricity is another pervasive network externality. Electric power follows Kirchhoff's law, as you may remember from Phys 1, so we don't know where the individual electrons are going but we know systemically what's going to happen. Electricity users are very sensitive to fluctuations in voltage—I have a set of expensive tube amps in my stereo at home; I'm really unhappy when they blow. And there are lots of computers containing lots of business records in lots of offices. So it's essential that the quality of the service is held constant; that is, the voltage fluctuations must be kept within tolerable levels.

The way this was traditionally dealt with was by having a monopoly; the monopoly solved the systemic problems; we regulated the monopoly. But a monopoly could charge a high price and be inefficient. Under regulation, it turns out it was *still* inefficient. Under the old regulatory system, we had balance—the system broke even; the price regulations ensured that. We had voluntary participation. The Supreme Court ruled in the Hope Natural Gas Company case of 1904 that a regulated firm was entitled to a fair return on its investment. So the utilities weren't coerced—Edison voluntarily sold electricity under regulation, and made money as a result. The argument

on the consumer side is a bit more subtle: nobody *forced* me to be part of the network. In theory, I could have always gone "off the grid" and put photovoltaic panels on my roof, or a windmill in my backyard. It wouldn't be cheap, but I could do it, and as electric bills spiral upward, a number of people are. Or I could have renounced my TV, microwave oven, air conditioner, computer, etc. and lived like the castaways on *Gilligan's Island*. It wasn't likely to happen, but nobody was stopping me. And, finally, the price caps made the system relatively strategy-proof. We had three out of four, and the downside was that we lost efficiency.

Deregulation was enacted under political constraints, so they went for four out of four. The theorem says you can't do it—unfortunately, nobody read the theorem in policyland. They intended to lower prices, so consumers would benefit and would join voluntarily; they were going to induce efficiency by relaxing producer price controls, so producers would join voluntarily, too; and the mechanism was set up to break even, so it had balance. And they relied on competition to make the mechanism strategy-proof.

So what happened? There's actually a sequence of markets. The so-called day-ahead market, for delivery tomorrow, is by the hour: 10 o'clock, 11 o'clock, 12 o'clock, and so on. The day-ahead market matches expected supply and demand. But say the next morning it turns out that the day is going to be hotter than forecast, and people are going to crank up their air conditioners. So there's the morning market, which is for same-day delivery in 15-minute intervals, to fine-tune supply and demand. And as the delivery deadline approaches, there are many more markets: a market for spinning reserves—people being paid to keep their generators running in case they're called upon; for nonspinning reserves—people who have their plants fired up, but the generators aren't turning; and so on. There's this whole hierarchy of markets based on how quickly a particular plant can be called on to produce. Then, at the very last moment, there's a market that forces supply to be equal to demand.

The later markets are run by the Independent System Operator, or ISO, a sort of quasigovernmental operation that pays whatever price has to be paid to maintain our constant voltage. The ISO's mandate is to keep the lights on at all costs. The ISO loses money, but it has to break even because the system has to be balanced. So it allocates its cost to the users, which in this case include the Big Three utilities: Southern California Edison, Pacific Gas & Electric, and San Diego Gas & Electric. But it doesn't instantaneously know who caused the excess demand, because it's a systemic problem, so the cost is shared via some rule.

In May 2000, for reasons that I'll talk about shortly, prices on the day-ahead market jumped through the roof. Consumer demand was up, so if the mechanism was strategy-proof, Edison, PG&E, and SDG&E should just have increased their demand and paid the higher price. But the cost they could pass on to the consumers was fixed by the retail price caps, so they didn't want to do that—they'd take a bath if they did. (As you know, they took a bath anyhow, but I'll get to that in a moment.) On the other hand, if they reduced their demand a little bit, it would push the unfilled consumer demand into those last-ditch markets where the ISO would have to cover it. The ISO divides its cost between the users, so the logic was this: I could buy an extra dollar's worth of electricity today, but if I don't, the ISO will buy it tomorrow. It will cost the ISO two bucks, but if the ISO divides that equally among us, two over three is less than a buck. Unfortunately, this only works if I'm the only person who does it. If two people do it, the cost becomes two people times two dollars divided by three people, which is more than a buck; if all three do it, everyone winds up paying two bucks. There's a strong incentive to be the first to act, even though the advantage you get is fleeting, because if you are honest you are guaranteed to lose unless everyone else is equally honest. (This particular scenario is a staple of game theory, and is called the Prisoner's Dilemma, because it was originally couched in terms of two cellmates given the opportunity to rat each other out in exchange for a lighter sentence.) So when the price went up, the declared demand—the amount the utilities said they wanted to buy—went down, and the excess was pushed into the residual markets. Suddenly the ISO, which was intended to do the transactions needed to suppress the last little fluctuations in the system, was buying 15 percent of the power. It was never meant to do that. And it was allocating costs in a way that was completely non-strategy-proof.

The supply shock—that price jump—was set up, again, because there was no balanced, strategy-proof mechanism. The miracle is that it took a year for the flaw to become apparent. Anyway, when supply was withdrawn, the same thing happened—the demand was forced into the last-minute markets, where you can charge almost anything and the ISO has to pay it. All at once, the scheduled maintenance time, the downtime when generators were removed from the system, roughly doubled. Maintenance outages are a matter of public record—whenever a generator is out for part of a particular day (the data don't track duration) the operator has to file a report. Generators were going on the fritz left and right—on some days we had 30 percent outages. And one study estimates that the producers' profits went up by $6 billion. Once again, you can't get four out of four. If we'd settled for three out of four, we never would have had this problem.

So the combination of giving the producers an incentive to withhold supply, magnified by an incentive for the buyers to withdraw demand

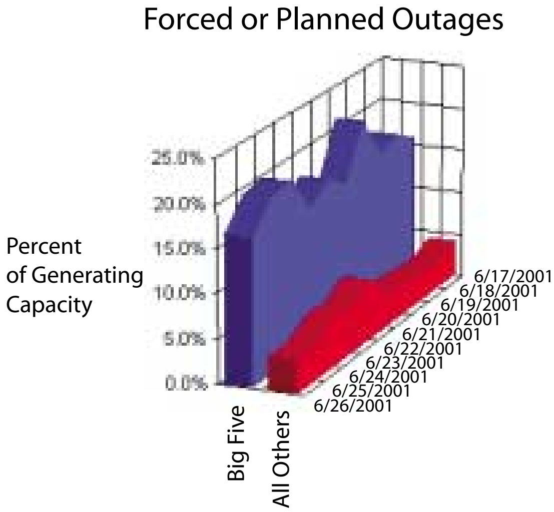

In the current market, when supply is reduced the sellers earn a higher price on every megawatt sold. Thus, when a generator goes off line, the sellers that own many generators make higher profits. The above data from the ISO shows a week's worth of generator outages reported by the so-called Big Five producers (AES, Duke, Dynergy, Mirant, and Reliant) compared with the smaller, independent producers.

PICTURE CREDITS: 28, 34–35 — Edison International; 30, 31 — Doug Smith

has proved to be disastrous. The system can't be balanced—it's running a giant deficit of about $15 billion, and the taxpayer has to pick up the bill. We have managed to achieve none out of four.

Since we can attain three out of four, there at least four possible solutions. We could renounce **E** (efficiency). That is, we can go back to a regulated monopoly, or we could introduce price caps. A lot of people are crying for that, because, well, things were bad in the old days but they weren't *this* bad! Or we could give up on **V** (voluntariness)—the state could seize the power plants through eminent domain, and force them to sell power to us at a fixed price. There have been a lot of calls for that, too. Alternatively, we could abandon **B** (balance). If we got rid of the balance requirement, we could assign long-term contracts for the delivery of a specified amount of power based on our best guesses for demand. We know we can engineer the awarding of these contracts in a strategy-proof manner, per Ted Groves's theorem I mentioned earlier. However, then the ISO will not always break even, and the taxpayer will have to foot the bill. And the imbalance could be large, in which case we're no better off than we are now.

My preferred solution is to design a better market; that is, we relax **S** (global strategy-proofness) and go for local strategy-proofness by giving people the correct incentives. In the previous market design, all the units of electricity that were bid for less than the market-clearing price—the price of the lowest unsuccessful seller (if you arranged all the bids from lowest to highest, the lowest unsuccessful seller would be the first seller whose bid was not taken)—were sold at the market-clearing price. We could stand that on its head, by breaking up the market into a set of smaller markets for each unit of capacity—per 100 megawatts, say. So there's a market for the first 100 megawatts, and the market sets a price. But then what we do is we award the sale—at the market-clearing price—to the generator who submitted the *lowest* bid. So it's to your advantage to bid a low price, because you'll get paid the highest price. Then there's another market for the next 100, and the process repeats. Now, if you try to withhold supply, you take yourself out of all the markets except for the last one, so your action benefits you only on the last 100 megawatts—unlike the exisiting situation, where you would affect the price of *all* the megawatts sold. And by taking yourself out of the previous markets, you lose all the business transacted therein; if demand is less than you predict and the market never gets to the 100 megawatts you're holding out, you're only hurting yourself. The more markets, the better this mechanism works—within computational reason, of course.

Actually, I think the best mechanism is to have a spot market with the no-**S** solution I've just described, coupled with the ISO using long-term contracts for power reserves that are awarded by what I call the Teacher's Pet method. Brownie points are given to producers who consistently have the lowest prices on the spot market, or who have the best reliability record—i.e., the fewest maintenance outages. Then if two bidders come in at the same price, the one with the most brownie points wins. Many government agencies already do this—the Defense Department, for example, puts performance incentives into its contracts, and bases future awards on the contractor's history of cost overruns and so forth. This system could also be used to reward whistle-blowing companies by giving them major brownie points in the next round of contracts. However, as I said before, when the ISO enters into long-term contracts, the balance requirement goes out the window. So I call this the no-**BS** solution.

These kinds of issues are going to get even bigger and more complex as our economy increasingly becomes a network of networks. So we really need to sit down now, and figure out the arcane details of how these markets work, in order to head off future missteps. □

Assistant Professor of Economics Simon J. Wilkie earned his B. Comm. from the University of New South Wales in 1984, and his MS and PhD from the University of Rochester in 1988 and 1990, all in economics. He was a member of the technical staff at Bellcore, where he specialized in game theory as applied to telecommunications markets, before coming to Caltech as an instructor in economics in 1994. He was made an assistant professor the following year. This article is adapted from a Seminar Day talk.