

Borenstein, Severin

**Article**

## Electricity pricing that reflects its real-time cost

NBER Reporter Online

**Provided in Cooperation with:**

National Bureau of Economic Research (NBER), Cambridge, Mass.

*Suggested Citation:* Borenstein, Severin (2009) : Electricity pricing that reflects its real-time cost, NBER Reporter Online, National Bureau of Economic Research (NBER), Cambridge, MA, Iss. 1, pp. 9-12

This Version is available at:

<https://hdl.handle.net/10419/61942>

**Standard-Nutzungsbedingungen:**

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

**Terms of use:**

*Documents in EconStor may be saved and copied for your personal and scholarly purposes.*

*You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.*

*If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.*

# Electricity Pricing that Reflects Its Real-Time Cost

Severin Borenstein\*

Long before U.S. electricity restructuring began in the 1990s there was a recognition that the marginal cost of producing electricity could change significantly hour to hour. Combined with the high cost of storing electricity, this meant that the true opportunity cost of consuming electricity also would vary constantly. For many decades economists have argued that retail electricity prices should fluctuate accordingly — this is known as real-time pricing (RTP) — but the technology to meter hourly consumption and to communicate fluctuating prices was quite costly.

In the last half of the twentieth century, the industry created a system meant to approximate RTP with standard technology: “time-of-use” prices that varied systematically by time of day and day of the week, usually with a higher price Monday through Friday during business hours and a lower price at all other times. The two prices (or sometimes three, with an added “shoulder” pricing period) were set months in advance, however, and did not change to reflect system demand/supply balance on a daily basis. Because of the cost of even this simple pricing and metering scheme, it was used only for large commercial and industrial customers.

In a regulatory environment, two additional factors worked against adop-

---

\*Severin Borenstein is a Research Associate in the NBER’s Industrial Organization and Environmental and Energy Economics Programs. He is also Director of the University of California Energy Institute and the E.T. Grether Professor of Business Economics and Public Policy, Haas School of Business, University of California, Berkeley. His profile appears later in this issue.

tion of RTP. Under regulation, the utility nearly always charges prices that are based on some notion of average cost, including the accounting amortization of long-term capital expenditures. Such an approach is targeted at cost recovery, not efficient pricing. Also, regulated utilities may be less likely to appreciate one of the main attractions of RTP, the effect it has in shaving demand peaks and reducing the need for capital investment. If regulators allow utilities to earn generous returns on investment, or if the utility management simply wants to grow the company, a pricing strategy that constrains new capital investment is unlikely to be popular with managers.

It is not that utilities did not understand or calculate their marginal cost. In fact, engineers tasked with minimizing production costs were constantly calculating “system lambda,” the value of the production constraint, which corresponds directly to economic marginal cost. They needed this information in order to choose among different production resources. The information was just not used on the consumption side.

As metering technology improved, a few utilities began to experiment with RTP. The pioneer and still a leader in this regard is Georgia Power, a company that was, and remains, a traditional regulated utility. GP introduced its first RTP program in 1991 for large industrial customers. By 2000, nearly one-third of its entire electricity demand was on RTP.

Wholesale electricity markets were deregulated in many parts of the United States in the late 1990s. The idea was that electricity generation could be a competitive industry with many generators vying to sell their output into a common

power market. The underlying economic model for this market, however, required that prices occasionally rise to well above the marginal cost of producing most units of output in order for firms to earn operating profits on infra-marginal units, operating profits that allowed the firm to cover its capital cost, at least in expectation. In the simple framework of a constant marginal cost of each generator up to its capacity, this meant that the market had to sometimes clear “on the demand side.” That is, high prices would occur at times of high demand or reduced supply, and those high prices would cause quantities demanded to decline until they were in line with system capacity. Such price-responsive demand would constrain prices from jumping too high, whether the tight market was caused by a true supply shortage or an artificial shortage caused by some firms exercising market power. What went largely unnoticed at the time was that the technology and market organization to enable RTP was not in place in any of the markets headed towards deregulation. My own pre-deregulation work with Jim Bushnell, which forecast market power problems in deregulated electricity markets, just assumed that there would be some degree of real-time price response.<sup>1</sup>

When deregulated markets launched in California, Pennsylvania-New Jersey-Maryland, and New England in the late 1990s, retail customers could choose among retail providers who were buying power out of the wholesale power market. Since electricity is a homogeneous good delivered over a common-carrier infrastructure of transmission and distribution wires, there was no ability to dif-

ferentiate the product sold. In nearly all cases, the final delivery and metering of usage was also left to the still-regulated utility that was providing transmission and distribution services. The centralization of the metering service meant that even if a retailer wanted to offer new time-varying retail pricing structures it was difficult to actually do so.

Still, because these markets were in a period of excess capacity, prices remained low and steady at first, with the primary complaint coming from producers who argued that prices were too low to justify new investment. That changed dramatically in summer 2000 with the onset of the California electricity crisis that brought extremely high prices and a few isolated blackouts when total supply from generation didn't keep up with demand.

The view of most economists who have studied the electricity crisis is that it resulted from a true scarcity in the wholesale market greatly exacerbated by the ability of a few sellers to exercise significant market power. This is supported by my own work with Jim Bushnell and Frank Wolak, and a paper by Paul Joskow and Ed Kahn.<sup>2</sup> Virtually all economists agree that the outcome was exacerbated by the inability of the demand side of the market to respond to real or artificial supply shortages. This realization prompted my research stream on real-time electricity pricing.

It was recognized during the crisis that RTP would lower prices during a time of supply shortage and would reduce the incentive of sellers to exercise market power by making demand more elastic, thus greatly reducing the wealth transfer from consumers to producers. The efficiency effects, however, were much less well understood. In a 2003 paper, Stephen Holland and I explored the short-run pricing and long-run investment inefficiencies that result when some or all customers face retail prices that do not vary with the wholesale market.<sup>3</sup> The theoretical analysis showed that the competitive equilibrium in a market without RTP could be quite inefficient. Not only would it fail to attain first-best

pricing because prices would not move with marginal cost, it would not even result in the second-best (that is, least inefficient) non-varying retail price. The result would be inefficient investment levels, even given the constant-price constraint. We also showed that incentives for adoption of RTP, if there is some cost to adopt, such as metering, could be too weak or too strong from a societal point of view. Essentially, this is because the RTP adopters change the price and levels of investment for the non-adopters as well, an externality that can be positive or negative.

In the same paper, we used simulations to examine how large the societal gains from switching to RTP are likely to be.<sup>4</sup> These simulations used realistic production cost parameters to analyze how the long-run equilibrium investment and pricing would change as more customers moved from a time-invariant pricing plan to RTP. The result was significant and at the same time sobering.

It was significant in that the potential gains from RTP were almost certainly many times greater than the estimated costs of implementing such a program. In addition, the gains were largest for the first tranche of customers moved to RTP. In fact, with reasonable elasticity assumptions, it is likely that one-half of the possible total surplus gain could result from putting only one-third of all demand on RTP. This was important because the cost of implementing RTP at the residential level may be substantially higher—because each household consumes fairly little yet has nearly the same metering and billing costs as a large industrial customer—so an RTP program is likely to start with large industrial and commercial customers.

The results were sobering because as exciting as the prospect of “getting prices right” may be to economists, the potential gains were likely to be only 5 percent or less of the energy bill. And energy is generally only about half of the entire electricity bill, the remainder being transmission, distribution, and customer administration costs. It still amounted to hundreds of millions of

dollars in California, but it wasn't going to fundamentally change the cost of supplying electricity. The reason for this is worth highlighting: in an electric system that must always stand ready to meet all demand at the retail price, the cost of a constant-price structure is the need to hold substantial capacity that is hardly ever used. But utilities optimize by building “peaker plants” for this purpose, capacity that has low capital cost and high operating cost. The social cost of holding idle capacity of this form turns out to be not as great as one might think. The analysis, however, does not capture some other potential benefits of RTP, including reduced vulnerability to supplier market power and greater resiliency in emergency situations, such as transmission outages, so the simulation estimates are only a piece of the gains.

While RTP holds potential for real efficiency gains, it is unfortunately often confused with energy efficiency programs that are designed to reduce overall consumption. RTP is even occasionally touted for having environmental benefits. While it might cause decreased consumption in some cases, there is no evidence that the effect on net would generally be in that direction. The effect of redistributing consumption from peak to off-peak periods could have positive or negative environmental effects. Stephen Holland and Erin Mansur's work on the environmental effect of RTP shows that in many parts of the country where coal provides baseload power, smoothing demand is likely to increase most pollutants, including greenhouse gases, by redistributing more production to coal-fired generation.<sup>5</sup> They find, however, that California, which relies less on coal, is an exception where demand smoothing from RTP is likely to benefit the environment.

The industry and public policy debate about RTP in the years following the California electricity crisis brought out a broad range of producer and customer concerns about RTP. Many were easy to address. For instance, managers at regulated utilities worry that RTP would make it more difficult for them

to be assured they could earn revenues that cover their costs, but RTP can actually match revenues to utility costs better than the time-invariant price model.<sup>6</sup> But two concerns, in particular, merited further empirical study.

First, some customers would be winners and others losers with a switch to RTP. Those who consume disproportionate quantities at the most expensive times are being subsidized under time-invariant pricing and may be worse off if they cannot adjust their consumption substantially under RTP. Most energy managers in industrial and commercial customers seemed to think that their bills on average would rise significantly under RTP (even though total system costs would fall with the switch). Second, even customers who thought their consumption pattern was no more expensive to supply than the typical customer were still worried that their bills could be much more volatile under RTP. Some economists dismiss such concerns about variance and risk management within companies, but the effect is very real on company budgets and performance reviews of the managers responsible for electricity consumption.

I was able to study both of these topics using a confidential dataset of hourly consumption for 1142 large industrial and commercial customers in Northern California over a four-year period. In the paper on wealth transfers from RTP, I combined the consumption patterns of these customers with simulated and actual wholesale prices to examine how big the likely transfer would be.<sup>7</sup> The results showed that the transfers were likely to be smaller than one might think. Starting from the simple time-of-use pricing these customers already faced (three different preset prices for peak/shoulder/off-peak periods), and assuming no change in consumption pattern, more than 95 percent of customers would be likely to see their bills rise or fall by less than 10 percent. I then looked at how much the losers might mitigate this impact by responding to the price variation. Over the plausible range of short-run price elasticities, the effect is fairly

modest: with no price-response, about 55 percent of these customers would see their bills rise under RTP (the winners are on average larger consumers than the losers), but with a short-run elasticity of -0.1, the share drops to 44 percent. The results made clear that extremely few customers would see disruptive changes in their electricity bills, but that there would be a significant number of small losers. The remainder of the paper discussed various strategies that attempt to compensate losers without distorting their marginal consumption incentives.

Using the same dataset, I studied bill volatility under RTP.<sup>8</sup> Bill volatility is caused by consumption volatility, price volatility, and the covariance of the two. In electricity, departures from average consumption quantity and average real-time price tend to be positively correlated, exacerbating the variance of bills. Using monthly billing periods, I calculated customer bills under time-invariant, time-of-use, and real-time pricing schemes. After adjusting for seasonal variation, which should be easy to anticipate, I found that the coefficient of variation of a customer's bill is on average nearly five times larger under RTP than under the time-of-use structure that they typically face. I then examine a simple hedging scheme in which the customer buys its expected consumption quantity (seasonally adjusted) for each hour of the month at the actuarially fair price for that period. This hedge incorporates no additional information that the customer is likely to have about its business activities during the month, so it is very likely that a customer could refine it further. Still, even this simple hedge eliminates about 90 percent of the excess bill volatility attributable to RTP, leaving bills about 30 percent more volatile than under time-of-use pricing. I also explored a somewhat more sophisticated strategy that utilizes over-hedging (buying forward more than 100 percent of expected quantity demanded) to compensate for the positive price/quantity correlation and showed that this approach can lower bill volatility to be about the same as under time-of-use pricing.

Over the last decade, real-time pricing has continued to attract attention and even some adoption. New programs have appeared in Chicago, New York, Florida, and elsewhere. For many years, resistance to adoption has rested on implementation costs, but technology advances have undermined those arguments. Nonetheless, as with many economically attractive ideas, public policy adoption requires first examining a number of issues, real and imagined, that fall outside the strict confines of economic efficiency. In my research, I have attempted to address both the questions of economic efficiency and the broader economic questions of risk and redistribution that are part of the policy process.

---

<sup>1</sup> S. Borenstein and J.B. Bushnell, "An Empirical Analysis of the Potential for Market Power in California's Electricity Industry," NBER Working Paper No. 6463, March 1998, and *Journal of Industrial Economics*, 47(3), 1999, pp. 285–323.

<sup>2</sup> S. Borenstein, J.B. Bushnell, and F.A. Wolak, "Diagnosing Market Power in California's Restructured Wholesale Electricity Market," NBER Working Paper No. 7868, September 2000, and *American Economic Review*, 92(5), 2002, pp. 1376–1405; and P. Joskow and E. Kahn, "A Quantitative Analysis of Pricing Behavior in California's Wholesale Electricity Market During Summer 2000," NBER Working Paper No. 8157, March 2001, and *Energy Journal*, 23(1), 2002, pp. 1–35.

<sup>3</sup> S. Borenstein and S. Holland, "On the Efficiency of Competitive Electricity Markets with Time-Invariant Retail Prices," NBER Working Paper No. 9922, August 2003, and *RAND Journal of Economics*, 36(3), 2005, pp. 469–93.

<sup>4</sup> A more extensive version of the simulation analysis was later published as S. Borenstein, "The Long-Run Efficiency of Real-Time Electricity Pricing," *Energy Journal*, 26(3), 2005, pp. 1–24.

<sup>5</sup> S. Holland and E.T. Mansur, "Is Real-Time Pricing Green? The Environmental

*Impacts of Electricity Demand Variance,” NBER Working Paper No. 13508, October 2007, and Review of Economics and Statistics, 90(3), 2008, pp. 550–61.*

<sup>6</sup> *I discussed many of the implementation issues in “Time-Varying Retail Electricity Prices: Theory and Practice,” in Griffin*

*and Puller, eds., Electricity Deregulation: Choices and Challenges, Chicago: University of Chicago Press, 2005.*

<sup>7</sup> *S. Borenstein, “Wealth Transfers from Implementing Real-Time Retail Electricity Pricing,” NBER Working Paper No. 11594, September 2005, and Energy*

*Journal, 28(2), 2007, pp. 131–49.*

<sup>8</sup> *S. Borenstein, “Customer Risk from Real-Time Retail Electricity Pricing: Bill Volatility and Hedgability,” NBER Working Paper No. 12524, September 2006, and Energy Journal, 28(2), 2007, pp. 111–30.*

## Re-evaluating Learning

Esther Duflo\*

Developing countries have rapidly increased access to primary school, but the quality of education has remained low. Many children are now in school, but they are hardly learning. In India, for example, a 2007 nationwide survey by Pratham<sup>1</sup>, a large education nonprofit, found that 97 percent of the of-age children are in primary school, but only 51 percent of third graders could read a simple first-grade paragraph, and only 33 percent could do simple subtraction. If developing countries are to attain meaningful universal primary education, they must improve the quality of education.

This is a formidable task: for starters, rising enrollment, unaccompanied by additional budget outlays, has increased pressure on available resources. Classes in the lower grades often are very large, and the children arrive with wide-ranging levels of preparedness. These large and heterogeneous classes can challenge pedagogy. The curricula, set nationally and often inherited in large part from the colonial period, are not adapted to local challenges and needs. Too often, they presuppose competencies that many of the first-generation learners do not

have. Besides these challenges, teachers face lax incentives, so teacher motivation is low: many teachers do not come to school and even those who do come do not always teach.

What can be done to improve education quality in developing countries? My recent research suggests some answers to this question. My approach has centered on using randomized evaluations to identify the causal effects of promising education programs.

In a randomized evaluation, from the program’s inception the researcher works in close collaboration with the practitioner. The program gets assigned randomly to part of the sample — the treatment group — which is compared to the rest, the comparison group. In recent years, there has been an explosion in research using randomized evaluations in development economics. Development economists have pioneered the use of research partnerships with non-governmental organizations (NGOs) or private companies. These partnerships often allow greater control over the research design and, increasingly often, input into the program design itself. Rachel Glennerster, Michael Kremer, and I<sup>2</sup> describe the various ways of incorporating random assignment in the evaluation design, and the practical challenges that go with it.

In “The Experimental Approach

to Development Economics”, Abhijit Banerjee and I<sup>3</sup> review the evolution of the use of randomization in development economics research. Much like earlier work in labor economics, health, and education, the experimental research in development economics started with concerns over the reliable identification of program effects in the face of complex and multiple channels of causality. The central difficulty that randomization seeks to address is selection bias. When program participants are not randomly selected, their outcomes may differ systematically from those of non-participants. This makes it difficult to attribute any differences observed between participants and non-participants to the program itself. For example, schools that receive better inputs also may differ systematically from the other schools in other ways, for example in pedagogy and teacher incentives. However, when the program is randomly assigned, these initial differences even out and selection bias disappears. Experiments allow researchers to vary one factor at a time by randomly assigning the program to part of the sample, and therefore they yield internally valid estimates of program effects.

Thus, in the mid-1990s, development economists started doing experiments to answer basic questions about the education production function:

\*Duflo is a Research Assistant in the NBER’s Program on Children and a professor of economics at MIT. Her profile appears later in this issue.