Note No. 87

Price Caps, Rate-of-Return Regulation, and the Cost of Capital

Ian Alexander and Timothy Irwin

This Note compares the effects of price cap and rate-of-return regulation on the risks borne by regulated utilities. It presents evidence that price cap regulation subjects firms to greater risks and therefore raises their cost of capital. This result has one clear implication: firms regulated by price caps must be permitted to earn higher returns. If they are not, they will be unable to attract new investment capital and the quality of their service will decline.

Price caps and rate-of-return regulation

There are two main approaches to preventing monopolistic infrastructure firms from charging excessively high prices: price cap regulation and rate-of-return regulation. The rate-of-return approach is used in Canada, Japan, and the United States, where regulatory agencies fix the rate of return that a utility can earn on its assets. They set the price the utility can charge so as to allow it to earn a specified rate of return—and no more. The regulated price can be adjusted upward if the utility starts making a lower rate of return, and it will be adjusted downward if the utility makes a higher rate.

Over the past decade or so, the price cap approach has become increasingly common internationally because it is thought to give firms stronger incentives to be efficient. Under this approach, the regulated price is adjusted each year by the rate of inflation plus or minus some predetermined amount and without regard to changes in the firm's profits. In the United Kingdom, for example, utilities are permitted to increase their prices by the change in the consumer price index plus or minus a specified amount. In gas and electricity, the price-setting rule is called $RPI - X,$ where RPI is the retail price index and X represents the expected annual gain in the utility's efficiency. In water,

the rule is $RPI + K$, where K represents both expected productivity gains and a permitted annual increase in the real price of water to allow for quality improvements (think of it as $RPI - X + Q$, where Q stands for the quality improvement). Since 1989, price caps have also been used in the United States to adjust the prices charged by the long-distance telephone company AT&T. In New Zealand, a price cap is used to adjust Telecom New Zealand's rental charge for a residential phone line. Price caps are also used in some developing countries. Malaysia, Mexico, and Peru, for example, use them for telecommunications, and Argentina uses them for gas and electricity as well.

In practice, price cap and rate-of-return regulation are less different than they might seem. First, a rule like $RPI - X$ considers only how prices should be changed from year to year; it doesn't tell a regulator how to set them in the first year. A regulator wanting to use price cap regulation for a new service would need to set the initial price in some way, and one obvious option is to consider the price the firm needs to charge to earn a satisfactory rate of return. Second, a price cap needs to be periodically reviewed: a regulator cannot reliably predict what changes in productivity will be possible in, say, ten years. In the United Kingdom, price caps typically are reviewed every five years. And during a review,

TABLE 1   AVERAGE INFRASTRUCTURE FIRM BETAS, BY COUNTRY, SECTOR, AND TYPE OF REGULATION, 1990–94

Country	Electricity Regulation	Beta	Gas Regulation	Beta	Combined gas and electricity Regulation	Beta	Water Regulation	Beta	Telecoms Regulation	Beta
Canada	—	—	—	—	ROR	0.25	—	—	ROR	0.31
Japan	ROR	0.43	—	—	—	—	—	—	ROR	0.62
Sweden	—	—	—	—	—	—	—	—	Price cap	0.50
United Kingdom	—	—	Price cap	0.84	—	—	Price cap	0.67	Price cap	0.87
United States	ROR	0.30	ROR	0.20	ROR	0.25	ROR	0.29	Price cap (AT&T)	0.72
									ROR (others)	0.52

— Not available or not applicable.

Note: The betas are asset betas that control for differences in debt-equity ratios between firms. ROR is rate-of-return regulation.

Source: Oxford Economic Research Associates, "Regulatory Structure and Risk: An International Comparison" (London, 1996).

the regulator naturally takes into account the regulated utility's rate of return. If it is high, the price cap is likely to be reduced; if it is low, the price cap may be relaxed.

But as long as price cap reviews are sufficiently infrequent (say, every five years), price cap and rate-of-return regulation should have different effects on regulated firms. In particular, a price cap subjects businesses to more risk. For example, under price cap regulation, if a firm's costs rise, its profits will fall because it cannot raise its prices to compensate for the cost increases—at least until the next price review, which may be several years away. Under rate-of-return regulation, however, the business would seek—and typically be granted within a year or so—a compensating price rise, so its profits would not change much. But if the firm's costs fall, price cap regulation is more advantageous to the firm than rate-of-return regulation, because it would retain more of the resulting benefits as profits. Thus, under rate-of-return regulation, consumers bear some of the risk that firms bear in price cap systems. This difference in impact means that firms subject to price cap regulation have a stronger incentive to lower their costs because they keep more of the cost savings than they would if they were subject to rate-of-return regulation. But the increased risk they bear tends to raise their cost of capital.

Empirical evidence on risk and the regulatory system

The risk that affects a firm's capital cost can be measured by a statistic called the firm's beta. Betas are used by investors worldwide and are an important factor in their decisionmaking. A firm's beta measures the extent to which the firm's returns vary relative to those of a diversified portfolio of equity holdings. It indicates whether an investor with a diversified portfolio would take on more risk by investing in a particular firm. The higher the beta, the bigger the increase in the riskiness of the investor's portfolio.

Several studies that compared the betas of British firms subject to price cap regulation with those of U.S. firms subject to rate-of-return regulation found that the U.S. firms have lower betas, as expected. But the results leave room for doubt because it is unclear whether it is the difference in regulation that's at work or something else, such as a difference in the degree of competition in the British and U.S. markets. But new empirical work done by Oxford Economic Research Associates for the World Bank tends to confirm the earlier conclusions. This study measured the betas of more than 100 infrastructure firms in many countries. Table 1 summarizes the results of the study, by country, for companies subject to price cap or rate-of-return regulation. (Some countries in the study have

been omitted from the table because they use discretionary regulatory regimes that do not follow a price cap or rate-of-return rule, or because the data were not comparable.) The results show that price cap regulation is associated with higher betas than rate-of-return regulation in Canada, Japan, and Sweden, as well as in the United Kingdom and the United States. Rate-of-return regulation is associated with betas ranging from as little as 0.2 in the U.S. gas industry to 0.62 in Japanese telecommunications, while price cap regulation is associated with betas ranging from 0.5 in Swedish telecommunications to 0.87 in British telecommunications. Overall, and as explained below, the differences in betas imply that firms subject to price cap regulation have to pay about an extra percentage point for their capital.

Why betas matter

To understand why betas matter, note that different firms face different costs of capital. Some firms must offer an expected rate of return of, say, 20 percent to attract investors, while others can get all the money they need by offering only 15 percent. Although the precise reasons for these discrepancies are not known with confidence, one critical factor is risk. Investors tend to be risk averse: other things equal, they prefer safer investments to risky ones. That means that firms have to compensate them for taking on more risk.

Investment risk, in the sense in which it is used here, relates only to bottom-line profits—the net impact on a firm's profits of all the separate risks facing the firm, such as operating risk, inflation risk, interest rate risk, foreign exchange risk, and political risk. Investment risk is not all downside. Risky projects are those with both a higher-than-average chance of turning out exceptionally badly and a higher-than-average chance of turning out well. Thus, when investors say they want to be compensated for taking on risk, what they mean is not just that they prefer an investment with a certain return of 10 percent to one that will probably make 10 percent but might make less. They mean that they prefer the safe 10 percent return to an investment offering, say, equal

chances of 5 percent and 15 percent returns. Investment risk, then, has to do with the variability of returns.

Much investment risk can be eliminated by diversification. To see why, consider a racetrack analogy. Suppose you have no information on how fast the horses can run. You could bet all your money on one horse, or you could bet a little on each horse in the race. The two strategies have about the same expected, or average, return: two people, each using one of the

Under rate-of-return systems, consumers bear some of the risks that firms bear in price cap systems

two strategies for, say, a thousand races, would probably end up with roughly the same amount of money. For any one race, however, the two strategies pose different degrees of risk. The strategy of betting on just one horse is riskier: you could do well, but you're more likely to lose everything you bet. But when you bet on every horse, you almost certainly will lose a little, because the racetrack has to make a profit.

As with betting on horses, investing in many firms eliminates much risk without significantly reducing the expected return. Thus, professional investors do not worry about the sort of risks that can be eliminated by portfolio diversification. But the risks of professional investment are different from those in racetrack betting. At the racetrack, you can eliminate almost all investment risk by betting on every horse. The same isn't true of investing in firms. Some years are good, and in these years, most firms do well. In other years, most firms do badly. So, on average, firms' returns tend to move in the same direction, and even if you've invested in every firm, the return on your portfolio is uncertain. This risk that remains after diversification is the risk that professional investors are most concerned about.

Professional investors are particularly interested in the likelihood that a firm's returns will move

with the returns on a completely diversified portfolio—that is, a portfolio that includes investments in enough firms so that further diversification would not significantly reduce risk. In one possible scenario, a firm's returns might be expected to vary in equal proportion to the diversified portfolio, so that, for example, when the returns on the portfolio increase by 10 percent, the returns on the investment also are expected to increase by 10 percent. In this scenario, beta equals 1, and the investment neither increases nor reduces the total riskiness of an investor's portfolio. As a result, investors will demand a moderate rate of return when investing in the firm, and the firm's cost of capital will be moderate.

In another scenario, a firm's returns might vary disproportionately with those of the diversified portfolio, so that a 10 percent increase in the portfolio's returns would be associated with, say, a 20 percent increase in the firm's returns, and a 10 percent decrease in the portfolio's returns with a 20 percent decline in the firm's. Here, beta equals 2. Because investing in such firms increases total risk, investors demand an above-average rate of return as compensation, and capital costs these firms more than it does the average firm.

In a third scenario, a firm's returns might vary less strongly with those of the diversified portfolio, with a 10 percent increase in the portfolio's returns associated on average with, say, a 5 percent increase in the firm's returns. Here, beta equals 0.5. Because investing in such firms reduces total risk, investors are willing to give up some return to invest in them. For these firms, the cost of capital is lower than average.

Betas and regulation revisited

Equipped with this measure of investment risk and the cost of capital, consider the returns available from investing in a utility subject to rate-of-return regulation. Because prices are adjusted each year to keep the rate of return roughly constant, investments in the firm are subject to little risk, particularly the market-related risk that investors worry about. If re-

turns in the market as a whole rise, the regulated utility's returns won't rise much (though they can rise a little in the period before the regulator requires a price cut). But if the market turns bad and returns fall, the utility's returns won't fall below the target set by the regulator for long. Thus, firms subject to rate-of-return regulation tend to have low betas and a lower-than-average cost of capital.

Price cap regulations don't have the same effect. Because in the short run the regulator sets no target rate of return, the regulated company's return can vary from period to period and is free to vary with the returns on the market. Even under price cap regulation, utility firms often have a fairly safe business, with returns that are affected less by economywide shocks than are those of other firms. As shown in table 1, their betas are still lower than 1, the average for all firms. But they are higher than the betas of firms subject to rate-of-return regulation. So investors will demand a higher return for investment in a firm subject to price cap regulation.

Conclusion

This does not imply that price caps are less desirable than rate-of-return regulation. It simply means that regulators need to take account of the effect of regulation on the cost the regulated firm has to pay investors for capital. Regulators using rate-of-return regulation can set the target rate of return lower than that earned by the average firm and still expect investors to be interested, because the returns are subject to less risk than those of an average firm. Regulators using price cap regulation need to give firms under their jurisdiction the opportunity to make somewhat higher returns, because those returns are riskier. If they don't, the firms will be unable to attract new investment capital, and the quality of their service will eventually suffer.

Ian Alexander, London Economics, London, and Timothy Irwin (tirwin@worldbank.org), Private Sector Development Department