3 The A-J Model under Uncertainty

3.1 Motivation

So far we have assumed that the firm knows exactly the profit it is able to earn with each combination of inputs. In actuality, a firm's profits are affected by many factors that are beyond its control and cannot be predicted perfectly. Several examples suggest the prevalence of this uncertainty. (1) Weather greatly affects the sales of energy utilities. When summers are mild, electric customers do not use their air conditioners as much as in hot summers, and as a result, the utility earns less profit. Similarly, mild winters reduce demand for gas and electricity for heating compared to more severe winters. Weather varies from year to year in unpredictable ways, such that the utility cannot fully know beforehand the profits that will result from its actions. Even over a period of years, cumulative profits may be higher or lower than expected due to long-term differences in actual weather from that predicted. (2) Firms must plan new capital facilities (e.g., power plants for energy utilities or fixed guideways for transit) years before the facilities will be operational. At the time the decision is made to construct new facilities, predictions are made concerning the building costs and the level of demand that will prevail when the facility is completed. Inevitably, actual costs and demand will differ from those predicted, and profits will not be as expected. (3) Operational costs for a firm vary from day to day due to such mundane factors as the number of employees unable to work on a particular day due to illness. When averaged over a period of time, costs might tend to stabilize. However, costs never become perfectly stable, and, more important, there may be no stabilizing tendency at all if shocks (such as strikes, unexpected inflation, and so on) occur more rapidly than the averaging over time can accommodate. Inevitably, the firm does not know, at the time of making decisions, exactly what its profits will be as a result of the actions it decides to take.

This uncertainty has important implications for the behavior of the firm and consequently for the design of appropriate regulation. The issue has been examined fairly extensively in relation to the Averch-Johnson model. Theorists have generalized the A-J model to incorporate the fact that the regulated firm is uncertain about the profits that result from its input choice. Peles and Stein (1976) show that some of the seemingly implausible implications of the standard A-J model (such as that the regulated firm increases its use of capital when the fair rate of return is lowered) do not occur when uncertainty is introduced. Das (1980) demonstrates, however, that the basic A-J result (namely, that the regulated firm uses an inefficiently high capital/labor ratio for its level of output) still holds when uncertainty is introduced. As such, these analyses strengthen the plausibility of the A-J model by removing problematical aspects while retaining its central message.

These investigations are also important, and perhaps more important, for their general implications beyond the confines of the A-J model. In particular, these analyses elucidate a fundamental concept regarding appropriate regulation in the face of uncertainty. They provide a case study of the fact that an asymmetric treatment of uncertainty—by which losses by the firm are treated differently by the regulator than extraordinary profits—leads to distortions in the firm's actions that operate against optimality.

This lesson has broad implications. There is often a tendency for regulators, in an attempt to protect the public, to treat extraordinary losses and gains differently. For example, if a firm makes a decision that later proves to have been wrong (e.g., to build a power plant that ends up being unneeded, or ends up costing much more than expected), there is a tendency to force the firm (namely, the shareholders) to bear part of the cost of the "mistake" rather than pass on the entire cost to the firm's customers. On the other hand, if the utility makes a decision that results in much greater profits than expected (e.g., negotiates long-term contracts for the supply of inputs just before the spot price of these inputs unexpectedly skyrockets), then there is a tendency for the gains to be passed on to customers, because allowing the firm to retain the extra profits would result in the firm earning more than a "fair" return. As the analysis in this chapter indicates, this asymmetry can actually induce the firm to make decisions in a way that ultimately works against the goals of the regulator and the welfare of customers.[1]

As a basis for comparison, section 3.2 describes the behavior of an *unregulated* firm facing uncertainty. Section 3.3 identifies the effect of rate-of-return regulation on the way uncertainty is perceived by the firm and the behavior of the regulated firm given this uncertainty. Several results are then demonstrated in section 3.4. The discussion throughout the chapter is motivated by the concepts of Peles/Stein and Das. However, our methods are different and less general, designed to allow the concepts to be visualized and to focus on the essential meanings that drive the analysis. In addition, we include some interesting results not discussed by these authors but implied by their work.

The findings can be summarized as follows for a firm in an uncertain world that is subjected to ROR regulation.

- Except in rare cases, the firm earns on average *less* than the allowed rate of return.
- Except in rare cases, the firm engages in pure waste.
- Contrary to the standard A-J result without uncertainty, the firm might utilize *less* capital than if it were not regulated. The reason is straightforward. If the risk the firm faces increases with the amount of capital that it employs, then ROR regulation provides the firm with two countervailing incentives. First, the firm has an incentive to *increase* capital so as to increase its allowed profit. Second, the firm has an incentive to *decrease* capital in order to reduce the asymmetrical risk it faces under ROR regulation. Either of these incentives could dominate, resulting in either an increase or decrease in capital.
- Contrary to the standard A-J result, lowering the fair rate of return can induce the firm to *reduce* its capital.
- The firm nevertheless chooses an inefficiently high capital/labor ratio for its level of output.

If the firm utilizes less capital than when unregulated, this last result implies that the firm also reduces its labor and output. The overall

1. If the "mistakes" of the firm are actually due to unreasonable behavior, requiring the firm to incur the subsequent loss is perhaps appropriate as a deterrent against such behavior. The point is simply that there is a tendency to apply a stronger standard of scrutiny to losses than to gains, forgetting that in an uncertain world both unexpected losses and unexpected gains are inevitable.

picture of ROR regulation under uncertainty is even more distressing than without uncertainty: in addition to utilizing an inefficient input mix, the firm probably wastes inputs and could easily produce less output than if it were not regulated.

3.2 Behavior of the Unregulated Firm under Uncertainty

Suppose that some intervening factor affects the profits that the firm is able to obtain at any input combination. For simplicity, let us suppose that the firm faces two possible events, called good luck and bad luck, where good luck means that the firm is able to earn greater profits at each input combination than under bad luck. For example, good luck might consist of prices for inputs dropping such that the cost of building a power plant is lower than expected, whereas bad luck is the opposite. Or, good luck might be an unexpected increase in demand such that the firm is able to earn greater profits at each input level, while bad luck is a drop in demand.[2]

With the exception of the last result of this chapter, the essential concepts concerning uncertainty can be described, and easily visualized, if we assume that there is only one input, which we call capital. At the time of choosing a level of capital (e.g., when deciding whether to build a new plant), the firm does not know whether good luck or bad luck will occur. Furthermore, the firm knows that it will not be able to adjust its capital after it has observed which type of luck has occurred. For example, after a new plant has been built, the sunk costs cannot be recovered if demand ends up being unexpectedly low such that the plant is not needed.

Because the firm does not know its exact profits at each level of capital, it calculates *expected* profits at each level of capital and chooses the level that provides the greatest expected profits.[3] For simplicity,

2. Note that good/bad luck simply refers to the effect of the event on the firm's profits. An event that constitutes bad luck for the firm might actually be socially beneficial. For example, mild weather reduces an energy utility's revenues and yet is beneficial for customers if the utility's rates are not raised to recoup its lost revenues.

3. We assume that the firm maximizes expected profits, though this need not be the case for all firms. A firm might maximize some other function of good- and bad-luck profits. For example, the firm might be more concerned about losing money than gaining, in which case it might maximize the weighted sum of expected profits and some factor that reflects the risk and harm of losing money. The tools used in this chapter can be generalized to allow for these other possibilities. However, the basic concepts regarding uncertainty are most apparent under expected profit maximization.

suppose that there is an equal probability that good luck and bad luck will prevail, such that expected profits are the average of the profits that would occur under good luck and those that would occur under bad luck.

Figure 3.1 illustrates the situation. There are two profit hills, one showing the profits the firm would earn at each level of capital if good luck occurred and another, lower one showing profits if bad luck occurred.[4] The average of these two profit hills is expected profits, denoted by a dotted line. The firm chooses the level of capital that gives it the highest expected profits, which is K_M.

The shape of the good- and bad-luck hills reflects the nature, or

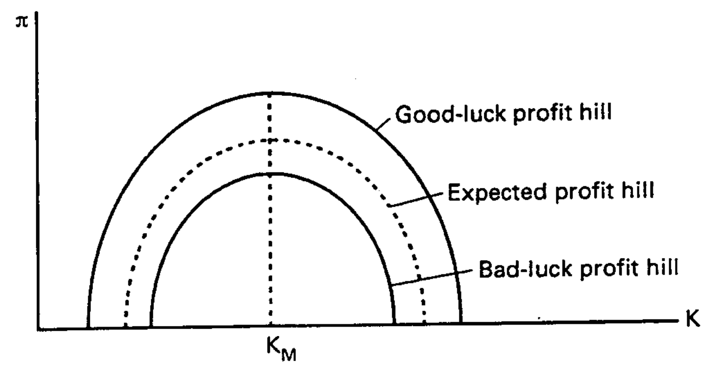

Figure 3.1
Unregulated firm facing independent risk

4. These profit hills, and the ensuing discussion, also have meaning in a two-input situation under a particular assumption about how each input is chosen. Suppose the inputs are capital and labor, and that capital is chosen before luck is revealed but that labor can be adjusted appropriately for the type of luck that occurs. This assumption is consistent with the idea that, for example, the sunk costs of a new plant cannot be recovered if demand ends up being insufficient to warrant the plant, but that the firm's labor force can be reduced or expanded in response to unanticipated events. Under this assumption, the good-luck profit hill gives the maximum profits that the firm can attain at each level of capital if good luck prevails and the firm adjusts its labor to be optimal given this good luck; and similarly for the bad-luck hill. All of the results in this chapter, with the important exception of the last one, are equivalently derived under this alternative interpretation of the profit hill.

All of the results, including the last one, can also be derived under the assumption that both labor and capital are chosen in advance and cannot be adjusted after the good or bad luck has occurred. However, the graphical methods in this chapter, while still suggestive, are not exactly accurate for the necessary demonstrations. The reason for this is that the firm would not necessarily choose beforehand the amount of labor that is optimal for the level of capital after the type of luck has been revealed. Consequently, the expected profit hill is not necessarily the average of the good- and bad-luck hills when expressed in the capital dimension only.

type, of uncertainty the firm faces. Two types of risk can be distinguished. (1) Independent risk refers to a situation in which the difference between good- and bad-luck profits is the same at all levels of capital. That is, the degree of risk is independent of the amount of capital employed. (2) Dependent risk occurs when the difference between good- and bad-luck profits varies with the amount of capital employed. That is, the degree of risk the firm faces depends on, and can be affected by, the firm's choice of capital.

In figure 3.1, risk is assumed to be independent: the good- and bad-luck profit hills are the same distance apart at each level of capital. As a result, the two hills reach their maxima at the same level of capital. The expected profit hill therefore reaches its maximum at the same level of capital as the good- and bad-luck hills. The firm chooses under uncertainty the same level of capital as it would if it knew with certainty which type of luck would occur.

Figure 3.2 illustrates a situation with dependent risk. As is usual with dependent risk, the difference between good- and bad-luck profits increases with capital. (As the firm uses more capital, it is usually more vulnerable to unexpected changes in input prices and demand because it has invested a larger amount of money in capital that cannot be altered after these unexpected changes have occurred.) The top of the expected profit hill occurs where the slope of this hill is zero. This point can be identified exactly by comparing the slopes of the good- and bad-luck hills. In particular, because expected profits is the average of good-luck and bad-luck profits, the slope of the expected profit hill is the average of the slopes of the good- and

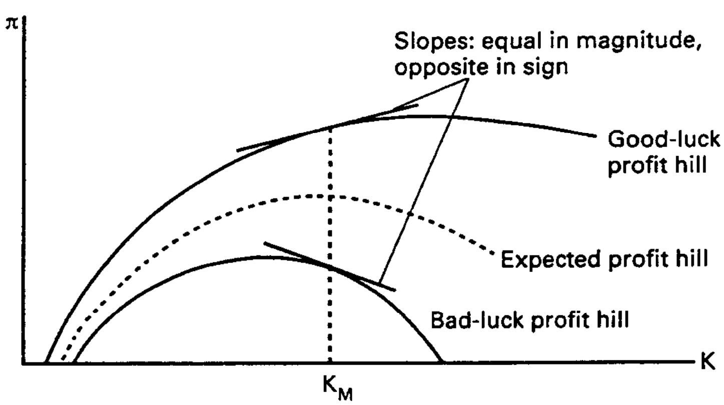

Figure 3.2
Unregulated firm facing dependent risk

bad-luck profit hills. The maximum of the expected profit hill occurs therefore at the level of capital at which the slope of the good-luck hill is equal in magnitude and opposite in sign to the slope of the bad-luck hill, such that the average of these two slopes is zero. This point is labeled K_M.[5] The firm in this case chooses more capital than it would if it knew with certainty that bad luck would prevail and less than if it knew good luck would occur.

3.3 Behavior of the Firm Facing Uncertainty under Rate-of-Return Regulation

Under ROR regulation, the firm is not allowed to earn more than a fair return on capital. This constraint interacts with the firm's good- and bad-luck profit hills, slicing off any parts that exceed the allowed amount. The firm's expected profits is the average of its sliced-off good-luck and bad-luck hills.

As figure 3.3 illustrates, under good luck, the firm would earn excessive profit over a range of capital levels; this part of the good-luck hill is sliced off by the constraint plane. Under bad luck, the firm, in this example, earns less than the allowed profit at all capital levels. The expected profit hill of the firm is the average of the bad-luck hill and the sliced-off good-luck hill; this average is denoted by the dotted line.[6] The maximum expected profits in this example occurs at capital K_R, at which the bad-luck hill slopes downward just as steeply (no more or less) as the constraint plane slopes upward, such that the average of their slopes is zero.

It is useful to identify the choice that the firm would make if it somehow knew the type of luck that would prevail. If the firm knew for sure that good luck would prevail, it would choose capital K_G, which provides the greatest profit under the sliced-off good-luck hill. If the firm knew for sure that bad luck would prevail, it would choose capital K_B, which is the top of the bad-luck hill. Because the firm does

5. The same relation holds with independent risk, though trivially. At the level of capital that maximizes expected profits, the slopes of the good-luck and bad-luck hills are each zero, such that their magnitudes are equal and their signs are irrelevant.

6. Note that this dotted line has a kink at the level of capital at which the constraint plane intersects the good-luck hill. For levels of capital immediately below this amount, expected profits is the average of allowed profits, which is the upward-sloping constraint plane, and bad-luck profits, which are downward sloping. For greater amounts of capital, expected profits is the average of the good-luck and bad-luck hills, both of which are downward sloping in this range.

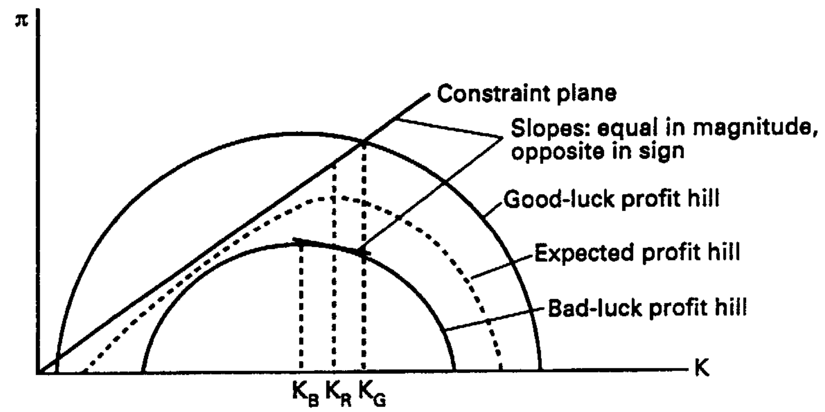

Figure 3.3
Example of regulated firm under uncertainty

not know the type of luck that will occur, it chooses K_R, which is between K_B and K_G.

In figure 3.3, the constraint plane slices the good-luck hill but not the bad-luck hill, and the chosen level of capital is strictly between K_B and K_G. Depending on the shapes of the good- and bad-luck hills and the allowed rate of return, the constraint plane may slice *both* the good- and bad-luck hills[7] and the chosen level of capital might equal K_B or K_G. Figure 3.4 illustrates these possibilities. In panel (a) the constraint plane slices both the good- and bad-luck hills. The chosen level of capital is identified the same as in figure 3.3, namely, where the negative slope of the bad-luck hill is the same in magnitude as the positive slope of the constraint plane. At this point, K_R is strictly between K_B and K_G. In panel (b) the constraint plane also intersects both hills. However, in this case, the bad-luck hill is steeper on the downward-sloping side than in panel (a). For all levels of capital beyond K_B, the bad-luck hill slopes down more steeply than the constraint plane slopes up. Expected profits therefore decrease beyond K_B. The firm chooses $K_R = K_B$, where expected profit is highest. In this situation, the regulated firm facing uncertainty chooses the same amount of capital as it would if it knew for sure that bad luck would prevail. Panel (c) depicts the opposite situation. Beyond K_B, the bad-luck hill slopes down *less* steeply than the constraint plane slopes up. Expected profit is highest at the level of capital at which the constraint plane inter-

7. For the regulation to be effective, the constraint plane must slice off part of the good-luck hill. As in chapter 1, the possibility that regulation has no effect is ignored.

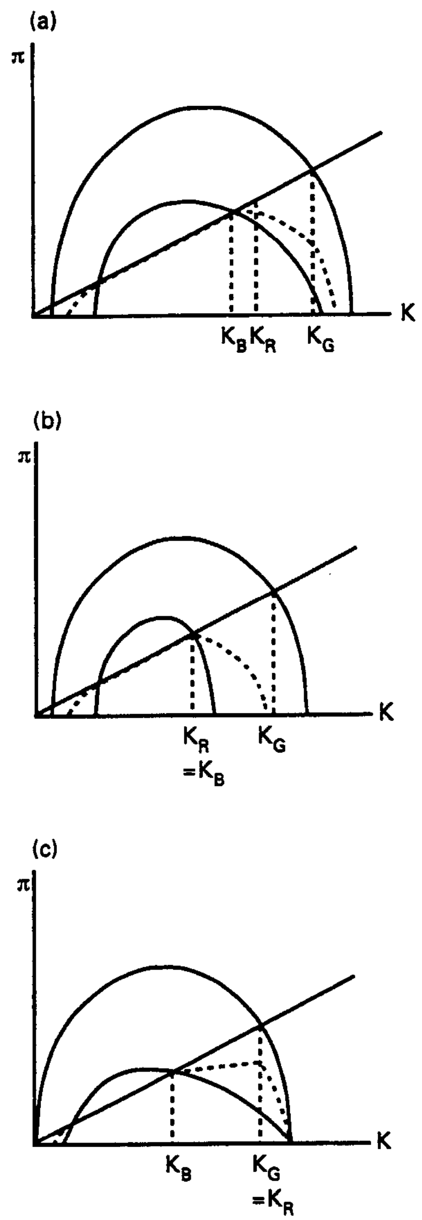

Figure 3.4
Other examples of regulated firm under uncertainty

sects the good-luck hill. This is the amount of capital the firm would choose if it knew with certainty that good luck would prevail.

Figures 3.3 and 3.4 indicate that the regulated firm under uncertainty chooses a level of capital at or between, but not beyond, the levels it would choose if it knew for sure which type of luck would prevail: $K_B \leq K_R \leq K_G$. The exact level of K_R within this range depends on the shapes of the good- and bad-luck hills and the slope of the constraint plane.

3.4 Results

Some results can now be demonstrated regarding the behavior of the regulated firm under uncertainty.

Result 1: If K_R *is strictly greater than* K_B *(that is, except in cases like panel (b) of figure 3.4), then the expected rate of return for the regulated firm is lower than the allowed rate of return. The regulated firm earns less profit on average than it is allowed to earn.*

In figure 3.5, the firm chooses K_R. If good luck prevails, the firm is capable of attaining profit π_G. However, because this profit exceeds the allowed amount, the firm is allowed to retain only π_A, the allowed profit. If bad luck prevails, the firm earns π_B. Expected profit, π_E, is the average of π_A and π_B, which is less than allowed profit π_A.

This phenomenon occurs whenever it is possible for the firm to earn *less* than the allowed rate of return but the regulator prevents

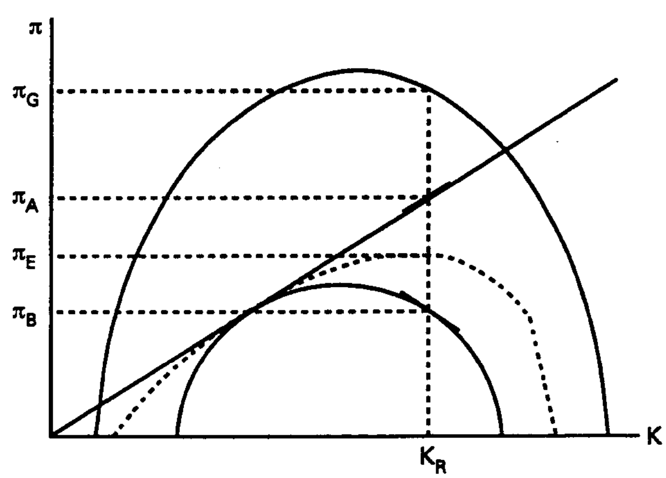

Figure 3.5
Regulated firm's expected profits are less than allowed profits

the firm from earning *more* than the allowed rate. The one case in which the firm's expected profits equal the allowed profits is when, as in panel (b) of figure 3.4, the firm earns exactly the allowed profit rate under both good and bad luck. Except in this rare situation, the firm on average earns less than the fair return.

There is a simple but often overlooked message in this result. If the regulator considers f to be the fair rate of return, the allowed rate should be set above f for the firm to have an expected profit of f. This concept is particularly meaningful in the context of demand fluctuations (due, for example, to weather). If the firm makes less than the allowed rate of return in "bad" years and yet is not allowed to make more than the fair return in "good" years, the firm's average return over time is less than the allowed return. If the regulator truly considers a certain rate of return to be fair, the allowed rate in each year must be set above this fair rate in order for the average rate that the firm earns over several years to end up being fair.

One way in which regulators have addressed this issue is to allow losses from one year (or, more precisely, the amount by which profits fall short of the allowed level in one year) to be made up in future years. That is, allowed profits in future years are raised until the losses are recouped. An accounting procedure is used to keep track of the cumulative sum of excess and deficient profits over time. When the sum is zero, the firm has earned the fair return over the period of time. This practice, and variants of it, are employed extensively, especially in relation to weather-related fluctuations in demand.

Result 2: If K_R *is strictly less than* K_G *(that is, except in cases like panel (c) of figure 3.4), then the regulated firm wastes if good luck prevails. On average, inputs are wasted.*

Figure 3.5 illustrates the situation. If good luck occurs, the firm can earn π_G if it operates efficiently. However, because it is only able to earn π_A, the firm must engage in some form of waste to reduce its profits by the amount $\pi_G - \pi_A$. This waste can take the form of producing less output than is maximally possible with the available capital.[8] This phenomenon of expected waste occurs whenever the firm, at its chosen capital level, is able to earn more than the allowed rate when luck is good. The one situation in which waste does not occur

8. If the firm uses labor along with capital and chooses its labor after knowing that good luck occurred, waste can be in the form of utilizing too much or too little labor relative to the cost-minimizing quantity.

is when, as in panel (c) of figure 3.4, the firm is able, if fully efficient, to earn only as much as the allowed rate when good luck occurs.

The relation of result 2 to the analogous result in chapter 1 provides insight into the source of waste in regulated firms. Recall that the standard A-J model implies that a regulated firm facing a fair rate of return that exceeds the cost of capital does not waste in the sense of producing less than is maximally possible with its inputs. The A-J model has often been criticized for this implication because it seems to contradict causal empiricism regarding the operations of regulated firms. The introduction of uncertainty into the A-J model provides information on a possible source of observed waste. In particular, waste can result from an asymmetric treatment of uncertainty. Waste occurs whenever the firm's allowed profit at its chosen capital level is less than the profit it is maximally able to earn with that capital. Thus waste can occur even if the firm is under some form of regulation other than ROR regulation, as long as the regulatory mechanism treats excess and deficient profits differently and the firm is maximally able to earn more than the allowed profit if good luck prevails.

If bad luck prevails, the firm does not waste. However, this fact has less importance than it might seem to. The explanations and graphs in this chapter are expressed in terms of two possibilities, good and bad luck. However, in the real world the possibilities are infinite. There is a continuum of profit hills, ranging from the very worst possible luck to the very best luck possible. For many of these (and perhaps all except the very worst luck) the firm wastes to some degree after its luck has been revealed.

Result 3: Under independent risk, the regulated firm utilizes more capital than the unregulated firm. However, under dependent risk, the comparison can go either way; in particular, the regulated firm might utilize less *capital than the unregulated firm.*

If risk is independent of the firm's choice of capital, the unregulated firm faces good- and bad-luck hills that reach their maxima at the same level of capital. The unregulated firm chooses this level of capital. At this level of capital, denoted as K_M in figure 3.6, the good- and bad-luck hills each has a slope of zero (such that the slope of the unregulated firm's expected profit hill is also zero).

When regulation is imposed, the top of the good-luck hill is sliced off; however, suppose for now that, as in figure 3.6, the top of the bad-luck hill is not sliced off. At the level of capital that the unregulated firm chooses, K_M, the slope of the regulated firm's expected profit

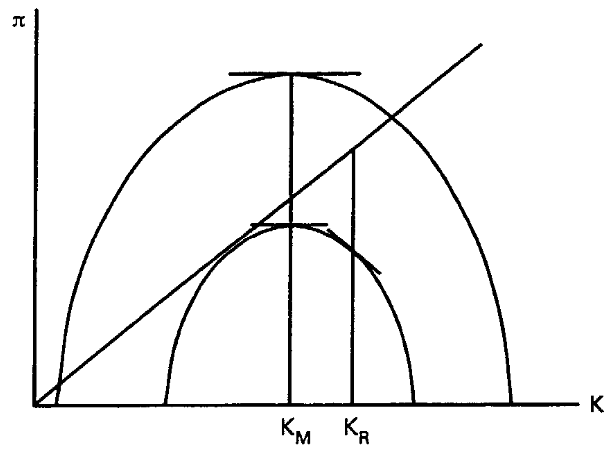

Figure 3.6
Regulated firm uses more capital than unregulated firm when risk is independent

hill is the average of the constraint's slope (which is positive) and the slope of the bad-luck hill (which is zero at K_M). Because the average of a positive number and zero is positive, the expected profit hill for the regulated firm is necessarily upward sloping at K_M, meaning that the regulated firm obtains greater expected profit at a higher level of capital. The firm therefore chooses greater capital, namely, K_R in the graph.

If the constraint plane slices off the top of the bad-luck hill as well as the good-luck hill, the same conclusion obtains. In this case, the regulated firm's expected profit hill near K_M is the constraint plane (see figure 3.4), which has a positive slope.

Consider now the possibility of dependent risk. In this case the regulated firm might actually choose *less* capital than if unregulated. Figure 3.7 illustrates such a situation. The unregulated firm chooses K_M, at which the slope of the good-luck hill is equal in magnitude and opposite in sign from that of the bad-luck hill (such that the slope of its expected profit hill is zero.) When the constraint is imposed, the slope of the regulated firm's expected profit hill at K_M becomes the average of the slopes of the constraint and the bad-luck hill. If, as in the figure, the bad-luck hill at K_M slopes down more steeply than the constraint plane slopes up, then the regulated firm's expected profit hill has a negative slope at K_M.[9] Because expected profits for the reg-

9. Another way to describe the situation is that the constraint plane slopes up less steeply than the good-luck hill at K_M. (Both slope upward, but the good-luck hill does

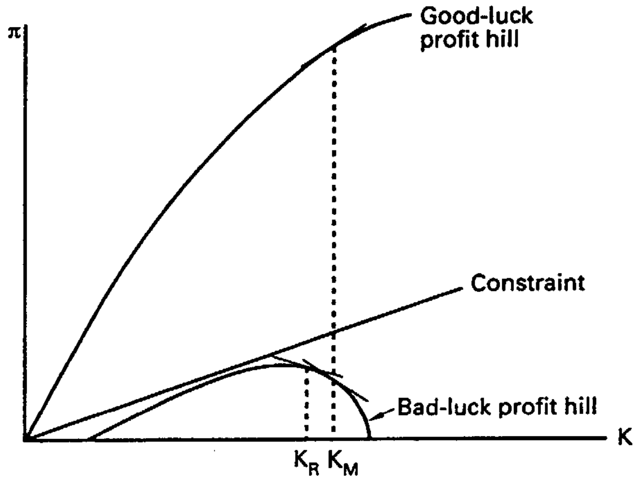

Figure 3.7
Regulated firm might use less capital than unregulated firm when risk is dependent

ulated firm are downward sloping at K_M, the firm obtains greater expected profits by utilizing less capital.

There is a clear reason for the firm facing dependent risk to reduce its capital when regulation is imposed. Without regulation, the firm might incur losses for reasons beyond its control, but it might also happen to obtain extraordinary profit. When regulation is imposed, the firm is still at risk of incurring the losses, but it is prevented from retaining any extraordinary profit that might occur, insofar as these exceed allowed profit. The larger the degree of uncertainty, the more the regulated firm is at risk of loss without a compensating chance of gain. As a result, the firm has an incentive to reduce its risk; and if risk is related to its use of capital, the firm has an incentive to reduce its capital. ROR regulation with dependent risk therefore provides the firm with two opposing incentives: to increase capital so as to increase allowed profits, and to decrease capital so as to decrease risk. Which of these two incentives dominates, and therefore whether the firm increases or decreases its capital when regulation is imposed, depends on the allowed rate of return and the relation of risk to capital.

Result 4: If K_R *is strictly between* K_B *and* K_G*, then lowering the allowed rate of return reduces the regulated firm's use of capital, under either independent or dependent risk.*

---

so more steeply than the constraint plane.) The good-luck and bad-luck slopes average to zero at K_M. Because the constraint has a smaller slope than the good-luck hill, the average of the constraint's slope with the bad-luck slope is necessarily negative.

The interesting aspect of this result is that it holds for independent risk, even though by result 3 the firm increases its use of capital when regulation is imposed. We therefore demonstrate the result with independent risk; the case of dependent risk is analogous.

In figure 3.8 the firm chooses capital K_0 under the original allowed rate of return. At this level of capital, the slope of the bad-luck hill is equal in magnitude and opposite in sign to the slope of the original constraint. When the allowed rate of return is reduced, the constraint plane rotates downward, becoming less steep. Because the slope of the bad luck hill at K_0 is equal in magnitude to the slope of the original constraint, it is greater in magnitude than the slope of the new constraint. Therefore, the average of the slopes of the bad-luck hill at K_0 and the new constraint is necessarily negative. Because the firm's expected profit hill with the new constraint is downward sloping at K_0, the firm obtains greater expected profits with a lower level of capital.

The result is essentially due to the fact that the expected profit hill is the average of the constraint and the bad-luck hill, such that reducing the slope of the constraint reduces the slope of the expected profit hill. At the point of zero slope originally, the slope becomes negative under a lower allowed rate.

If K_R is originally equal to either K_B or K_G (that is, the level of capital that the firm would choose if it knew for certain that bad luck or good luck, respectively, would occur), then lowering the allowed rate of

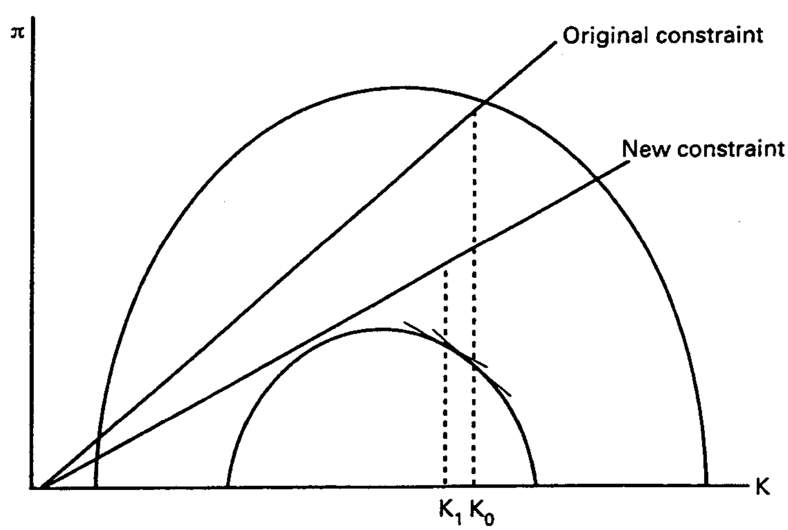

Figure 3.8
Regulated firm uses less capital when fair rate is reduced

return does not decrease the use of capital. In panels (b) and (c) of figure 3.4, lowering the constraint induces the firm to use more capital.

Result 4 is important because, like result 2, it reverses a widely criticized implication of the standard A-J model. Without uncertainty, the A-J model suggests that lowering the allowed rate of return necessarily increases the firm's use of capital, which is contrary to the concept that capital leaves industries whose profit rates fall.[10]

Results 2 and 3 are descriptive rather than normative, in that they describe how the regulated firm behaves in its choice of capital but do not indicate, directly at least, whether the level of capital is efficient. The following result suggests that the basic A-J effect, namely, that the firm uses an inefficient input mix, holds when uncertainty is introduced even though, as we have seen, other implications of the standard A-J model do not.

Result 5: The regulated firm's capital/labor ratio is inefficiently high.

For this result, it is necessary to include labor in the analysis of the firm's behavior. Assume that the firm chooses labor at the same time as capital; that is, assume that the firm chooses the levels of capital and labor without knowing whether good or bad luck will occur and cannot adjust these levels after its luck has been revealed.

I have not been able to devise a graphical method for demonstrating this result. However, in our situation involving only good or bad luck, a mathematical demonstration turns out to be simple and informative.

Assume that the constraint plane slices off all relevant portions of the good-luck profit hill but does not intersect the bad-luck hill, such that expected profit is the average of the constraint and the bad-luck hill. This is the two-input analog of figure 3.3; other situations can be examined analogously. Let $RB(K,L)$ be the largest revenues that the firm is able to attain under bad luck with inputs K and L. Given wage rate w, interest rate r, and allowed rate of return f, the firm's expected profit is

10. This criticism of the A-J model is based on an idea that in reality the supply of capital available to the firm depends on the firm's allowed rate of return, while in the A-J model the firm is assumed to be able to obtain as much capital as it wants at a given interest rate r. The direct way to approach this criticism is to generalize the A-J model to allow for a variable price of capital, with r being a function of the allowed rate of return and perhaps the amount of capital purchased by the firm. The analysis of uncertainty, however, indicates that the expected direction of effect can be obtained even without the generalization to variable r.

$$EP = (1/2)[(RB(K,L) - wL - rK)] + (1/2)[(f - r)K].$$

The first term in brackets is the profit the firm obtains under bad luck. Because we assume that the firm under bad luck cannot feasibly earn as much profit as it is allowed, the profit the firm attains is simply the maximum revenues minus the cost of the inputs. The second term in brackets is the profit under good luck. Because feasible profit is assumed to exceed allowed profit under good luck, the firm keeps only the allowed profit $(f - r)K$. Expected profit is the average of profits under good and bad luck, with a probability of 1/2 for each.

The firm chooses the K and L that maximize expected profit. Taking the partial derivative of EP with respect to capital:

$$\delta EP/\delta K = (1/2)[(RB_K - r) + (f - r)] = 0,$$

or,

$$RB_K = r - (f - r).$$

where RB_K is the marginal revenue product of capital. This marginal revenue product is the extra revenue the firm obtains when it utilizes an extra unit of capital and sells the extra output that it is able to produce with this extra capital. As such it can be expressed as the marginal product of capital times the marginal revenue of the firm's output: $RB_K = MP_K \cdot MR$. Substituting into the above equation:

$$MP_K \cdot MR = r - (f - r). \tag{3.1}$$

Taking the partial derivative of EP with respect to labor:

$$\delta EP/\delta L = (1/2)\ (RB_L - w) = 0,$$

or, because $RB_L = MP_L \cdot MR$,

$$MP_L \cdot MR = w. \tag{3.2}$$

The ratio of equation (3.1) to (3.2) is:

$$\frac{MP_K \cdot MR}{MP_L \cdot MR} = \frac{r - (f - r)}{w}$$

or

$$MP_K/MP_L = (r/w) - (f - r)/w.$$

If $f > r$, the term being subtracted is positive, such that the ratio of marginal products is *less* than the ratio of input prices. At the cost-minimizing input mix, the ratio of marginal products (the slope of the isoquant) is *equal* to the ratio of input prices (the slope of the isocost). The firm is therefore using an inefficient input mix. Furthermore, because the ratio of marginal products is *less* than the ratio of prices, costs would be lower with less capital and more labor.[11] That is, the firm chooses an inefficiently high capital/labor ratio.[12]

Results 2, 3, and 5 combine to paint an even more distressing picture of ROR regulation than obtains when uncertainty is not considered. When risk is dependent on the level of the firm's operation (as it usually is), then regulating the firm can induce it to reduce its level of capital. At this lower level of capital, the firm uses an inefficiently high capital/labor ratio, which implies that its labor and output are also reduced. Except under rare conditions, the firm also wastes on average. Stated succinctly, ROR regulation in an uncertain world can induce the firm to reduce its output, engage in pure waste, and purchase an inefficient input mix.

The forces driving these results are important to distinguish. By

---

11. This fact is elaborated in section 1.5. Assuming diminishing marginal products, the use of less capital increases the marginal product of capital, and the use of more labor decreases its marginal product. The ratio of marginal products therefore increases, rising toward the ratio of input prices as required for cost minimization.

12. The question naturally arises with regard to result 5: For what level of output is the input mix inefficient? For the analysis of result 5 to be fully consistent and meaningful, we can assume that the output of the firm is the same under both good and bad luck. Luck simply determines the *price* at which the firm can sell this output. (Under good luck, the firm can charge a higher price than under bad luck.) The difference between the good-luck and bad-luck profit hills is therefore attributable to this price difference. In this context, result 5 says that the firm chooses an inefficiently high capital/labor ratio for its level of output, which is the same under both good and bad luck.

Under this interpretation of the good- and bad-luck hills, the constraint on profits prevents the firm from raising price as high as it would otherwise be able to under good luck. Because the regulated price is below the highest price at which the firm could sell its output, demand for the firm's output exceeds the quantity produced when good luck occurs (assuming, as usual, downward-sloping demand). This is perhaps somewhat consistent with reality, in which, for example, an electric utility is not able to meet demand during extremely hot days when most air conditioners are on and running continuously. However, the concept that output is constant is contrary to some of the interpretations provided for other results, particularly result 2 regarding waste. If output is constant and only price differences determine the difference between the good- and bad-luck hills, the firm does not waste in the sense of producing less output than possible. Rather the inefficiency takes the form of excess demand. When price is not used to allocate available supply and excess demand arises as a consequence, customers who are most willing to pay for the output are not necessarily the customers who receive the output. This misallocation incurs a cost on society, which can be considered a form of waste. Consequently, result 2 can be considered to hold under the assumptions for result 5, but with the waste taking a different form.

tying allowed profits to the firm's use of capital, ROR regulation provides the firm with an incentive to substitute capital for labor. By treating the consequences of uncertain events asymmetrically, the firm is induced to reduce its risk, even if doing so means reducing output (by purchasing fewer inputs and/or wasting the inputs that it purchases). This reaction to uncertainty, though examined in the context of ROR regulation, can be expected to occur, in some form, under any regulatory mechanisms that let the firm incur losses (or less than allowed profits) due to chance events but do not allow the firm to retain excess profits when these are also due to chance events. By applying a stricter standard of review to windfall profits than to unexpected losses, the regulator, while trying to serve the interests of the public, actually induces behavior that is contrary to its own goals.