Costing Definitions and Concepts

Explanation of the term "economic cost," as contrasted with embedded cost.

The term "cost" is used in different contexts (and by different individuals) with different meanings. It is therefore useful to distinguish the accountant's use of the term from the economist's use.
Accountants are concerned primarily with the proper recording and measuring of historical costs based upon a uniform set of rules. They have developed a comprehensive system of recording and reporting data about costs, which is used by managers, investors, regulators, and economists in carrying out their respective jobs. The data, recorded in the books and records of a firm, are referred to as "accounting" or "embedded" costs. Accountants have also developed various "cost accounting" rules concerning how costs should be allocated to various categories.
Economists, on the other hand, have developed a comprehensive set of theories concerning cost, which they use to describe, explain, and predict the behavior of firms and individuals (e.g., consumers). The field of economics thus provides the underlying theory of costs, while accounting generally supplies most of the data that allow this theory to be applied in practice.
While embedded costs--the accountant's measure of cost--are quite practical, readily available, and fairly consistent from firm to firm, the economist's idea of cost is more useful in analyzing the critical decisions made by management and government.
In order to develop an appropriate costing methodology for a telecommunications service, it is important to understand both the underlying economic theory (and associated terminology) of cost, and the accountant's practical measures of cost (which do not directly correspond to elements of the theory).

Economics recognizes a variety of different types of cost:

Some of the most fundamental and important types of cost are total cost, variable cost, fixed cost, average cost, direct cost, joint cost, common cost, sunk cost, marginal cost, incremental cost, embedded cost, and fully allocated cost. Each term is applied to a separate and distinct concept; all but the last two are integral parts of economic theory.
Total cost is the sum of all costs incurred by the firm to produce any given level of output--that is, the sum of the firm's variable and fixed costs.
A fixed cost is invariant with the level of production, thereby not changing in the short run. Such costs must be paid regardless of how much the firm produces, or whether it produces at all, as long as the firm does not withdraw the factors of production entirely from the market by, for instance, selling its assets.
A variable cost changes directly (but not necessarily proportionately) with the level of production. The sum of the firm's fixed and variable costs equals its total cost of production.
Average total cost is the total cost of producing a given quantity of output, divided by the total number of units produced.
Average fixed cost is the summation of all fixed costs of producing a given quantity of output, divided by the total number of units produced.
Average variable cost is the summation of all variable costs of producing a given

quantity of output, divided by the total number of units produced.

Sunk costs have already been incurred, are considered irretrievable, and are thus irrelevant to current decisions, because they cannot be avoided regardless of the course of action selected. Although sunk costs are a major component of the total costs recorded on a firm's books, they are excluded in the context of economic costs. Hence, sunk costs have no impact on total, average or marginal costs--when these economic concepts are properly applied.

Marginal cost is the change in total cost resulting from an extremely small change in output. In mathematical terms, marginal cost is the first derivative of the total cost function with respect to output, assuming the cost function is continuous and smooth. Unlike average cost, marginal cost is not influenced by fixed costs. Marginal cost is affected only by variable costs. In practical applications, the cost function is not necessarily smooth or continuous. Accordingly, it is sometimes necessary to estimate the rate of change in cost over a discrete interval, using smoothing or averaging techniques. Nevertheless, when properly applied, marginal cost focuses on the effect of very small changes in output occurring at the point in the total cost curve where the firm is operating and decisions are being made.

Incremental cost is the change in total cost associated with a specified increase or decrease in output. Typically, incremental costs are reported on a per unit basis, and thus the change in total cost is divided by the number of units specified. In mathematical terms, incremental cost equals total cost assuming the increment is produced, minus total cost assuming the increment is not produced, (divided by the number of units in the increment). While marginal cost is always calculated with reference to a very small incremental change in output, the more generic term incremental cost can also be used in contexts where a very large output change is contemplated.

In firms that produce a variety of different products or services (multi-product firms) some additional distinctions are important. Each of these concepts can apply to total and average costs, including both fixed and variable components.

A direct cost can be specifically attributed to the production of an individual service or product, without requiring the use of allocations to separate it from costs incurred in the production of other services or products.

Common costs are incurred when production processes yield two or more outputs. They are often common to the entire output of the firm but can be common to just some of the outputs produced by the firm. An increase in production of any one good will tend to increase the level of common costs; however, the increase will not necessarily be proportional. The costs of producing several products within a single firm may be less than the sum of the analogous costs that would be incurred if each of the products were produced separately.

A joint cost is a specific type of common cost--one incurred when production processes yieldtwo or more outputs in fixed proportions. A classic example arises in the joint production of leather and beef. Although cattle feed is a necessary input for the production of both gloves and hamburgers, there is no economically meaningful way to separate out the feed costs that are required to produce each. If the quantity of leather and beef is reduced, there will be a savings in the amount of cattle feeding costs, but it is impossible to say how much of this change in cost results from the change in the quantity of leather and how much from the change in the quantity of beef.

An allocated cost is a joint or common cost that has been divided among the firm's different customers or products, in accordance with a particular formula or the judgments of a cost analyst.
Fully allocated costs are the summation of direct and allocated costs for a customer, customer class, product, or product group, developed in a cost study in which none of the firm's joint and common costs are left unallocated. Fully allocated costs are often referred to as fully distributed costs.

The relevance of this long list of cost definitions

Often, when significantly different cost estimates are presented to a regulatory commission, these differences can be traced in large part to fundamentally different definitions of cost; perhaps one analyst has estimated average cost while the other has estimated marginal cost (though both may appear on the surface to be estimates of the same thing). In understanding the differences between various cost estimates, it is also helpful to appreciate the theoretical distinctions between specific types of cost, such as joint, common, fixed, variable, and sunk costs.

Explanation of the difference between embedded direct and fully allocated embedded cost studies

These studies are quite similar in some regards. They both reflect accounting costs (i.e., the costs recorded on the books of the utility), rather than economic costs. They typically focus on broad categories of service rather than on individual services. For instance, the study might show the cost of local, intrastate toll, interstate access, private line, and miscellaneous services. Generally, it would not show the cost of individual products or services within those broad categories (e.g., a specific optional toll calling plan) or of specific pricing elements (e.g., daytime message toll calls spanning 25 miles).
The procedures used in developing both embedded direct and fully allocated studies are also quite similar. In both cases the analyst works with costs from the firm's accounting records and attempts to attribute these costs to specific categories of service.
The distinguishing feature of these two types of studies is their respective treatment of joint and common costs. Because of the savings that arise when multiple goods are produced collectively, joint and common costs cannot be meaningfully attributed to any single product or market. Differences in the way joint and common costs are analyzed largely explains the difference between embedded direct and fully allocated costs (as well as differences between particular fully allocated cost studies).
Properly conducted, an embedded direct study will assign only those costs that can be directly traced to a particular service category. Joint and common costs will be left unassigned, typically as one or more lump sum amounts. In fully allocated studies, however, no costs are left unassigned. The joint and common costs of production are allocated to the various categories of service under study, using various formulas that reflect relative usage or other factors.
The most widespread embedded direct study used in telecommunications was previously called the Embedded Direct Analysis (EDA). It was developed by AT&T and conducted annually by most Bell operating companies (BOCs) prior to their divestiture from AT&T. The EDA showed the direct embedded costs associated with such categories as local exchange, intrastate and interstate private line, intrastate and interstate toll, and various

miscellaneous supplemental services offered by the local BOCs. It was often presented in rate design proceedings to show the cost/rate relationships among these major service groups. Some of the former AT&T subsidiaries have apparently discontinued performing embedded direct studies; others have modified the EDA and given it a new name, such as the Cost Accounting Allocation System (CAAS).

Explanation of joint and common costs

A firm that produces a single product sold in a single market incurs only direct costs. These include capital costs (cost of money, depreciation, income taxes) and all expenses exclusively attributable to a specific product or service. However, when the firm is engaged in producing multiple products or serving multiple markets, it normally also incurs joint and/or common costs.

Joint costs are incurred when production facilities simultaneously serve two or more markets (or produce two or more products) in fixed proportions. Because proportions are fixed, it is impossible for the firm to increase or decrease the amount of output for one market without changing in the same proportion and in the same direction the output or capacity available for another market. Consequently, joint costs vary in proportion to the total available output of the joint production process, not the output of the individual joint products. Due to this interdependence between different products and markets, joint costs pose some special problems in the economic theory of pricing.

Common costs fall somewhere between direct and joint costs: they are not directly attributable to a single service, yet they vary to some degree with the quantity of production of each service. Typical examples of common costs include salaries and other costs of the firm's upper level executives, regulatory and legal expenses, and audit expense. All of these examples vary with output, but it can be exceedingly difficult to discern the extent to which they vary with the volume of production of any particular product or service.

'Local loop': joint or common cost?

In the telecommunications industry, the cost of the "subscriber loop" is a joint cost required for the provision of at least three different services: local exchange service, intrastate long-distance service, and interstate long-distance service. Since the installation of an additional subscriber loop increases the capacity available for placing and receiving all three types of calls, the telephone company cannot increase the capacity for local calls without concurrently increasing the capacity for toll calls. In this sense, it clearly fits within the definition of joint costs--since access capacity is simultaneously expanded for multiple services in fixed proportions (one more line is available in each case). Only if there is congestion at a particular time is there any tradeoff between use of the local loop for the different purposes.

Economic theory demonstrates that there is no inherently correct method of allocating joint costs among the various joint products. Purchasers of each of the joint products will bear some share of the joint costs, in relative proportions that are determined by the relative strength of demand in the various markets, rather than by some arbitrary allocation formula. In other words, cost recovery does not depend upon relative usage, nor does it depend upon any particular allocation scheme. Rather, it depends upon the interplay of supply and demand in the various markets involved.

How joint and common costs are recovered in competitive markets

To the extent common costs vary with output, they are recovered in the same manner as direct costs--they directly affect the marginal cost of producing each service, and thus directly influence prices. (In competitive markets, prices tend to equilibrate towards marginal cost). Joint costs, on the other hand, have no impact on marginal cost, and thus do not directly determine prices in competitive markets. In competitive markets joint costs are never recovered entirely from consumers of one of the joint products, to the exclusion of the others; rather, the costs are shared by both groups of consumers, with the respective proportions depending upon the relative strength of demand. The stronger the demand for a particular joint product, the greater the share of joint costs that will be borne by that product.
The joint costs of the loop capacity that serves both the toll and local markets. In fact, because the demand for toll service is stronger than the demand for local service on a per-minute basis, (reflecting the inherently greater value of communicating over much longer distances), a competitive market result would reflect a larger contribution (per minute) from toll customers than from local customers. The telecommunications companies disagree. Their usual view of joint loop costs it that they should be recovered entirely from local service rates.

The treatment of joint and common costs in the various types of cost studies

In the embedded direct study, only direct costs are attributed to the three categories. A comparison of these direct costs with the revenues generated within each category shows the extent to which each service is covering its direct costs and providing a contribution toward the firm's joint and common costs.
In the fully allocated study, all of the firm's costs are assigned to one of the three service categories. Usually, these studies show the investment and expense components of total cost separately so that a rate of return can be calculated for each service category. Joint and common costs are allocated, using some plausible procedure; for instance, these costs might be spread in proportion to revenues.
In a marginal cost study, joint costs should be excluded (or would mathematically have no impact on the result), because by definition they are fixed with respect to output of any one product or service. Since joint costs do not vary with output, they do not affect the marginal cost. Common costs, on the other hand, can potentially vary with output, and thus may have an impact on marginal cost. If one can determine the extent to which the common costs directly increase with an increase in production of a good or service, this would be one element of the marginal cost of producing the good or service in question. For simplicity, common costs are sometimes excluded from marginal cost estimates, with the understanding that the estimate will be biased downward, since common costs are excluded.

The handling of joint cost in a fully allocated cost study

The requirement in a fully allocated cost-of-service study that all costs be allocated, regardless of how ambiguous the causal relationship with the service in question, produces results that are defined by the particular allocation methodology selected, rather than by established economic costing principles. Thus, any number of widely different

estimates of "cost" could be produced for a given service category, merely by changing the allocation procedure. One study might show a particular service category earning an above-average rate of return, while another study of the same company might show a negative return for that category. The allocation scheme is pivotal. These allocation decisions are highly judgmental and (not surprisingly) controversial in regulatory proceedings where fully allocated studies are introduced--particularly where the joint costs are a very substantial fraction of the firms total costs.

For example, subscriber loop costs are joint costs, which constitute an extremely large and important component of a local exchange carriers cost structure. If TSIC or ISIC costing principles are strictly applied, most loop costs will be excluded from the cost estimates for all of the various services--since the loop costs do not generally vary with output of any single service. Rather, they vary with output of the entire family of services that require use of the loop.

For instance, as a carrier expands output, by adding additional customers, neighborhoods, or cities to its scope of operation, it will incur additional loop costs. However, it will also gain additional revenues in a wide variety of different services--including local, toll, access, and custom calling. By looking at the entire family of services that use the loop, one can determine whether the additional loop costs are adequately offset by additional revenues generated by these loops. However, if one focuses on one particular service (e.g. local exchange), a meaningful comparison cannot be directly made.

One would not expect any one service to produce enough revenues to recover the entire cost of the loop, nor would it be important for this to occur, since other revenue streams will also be enlarged as the company expands. Similarly, cattle breeders are not concerned with whether beef prices are high enough to recover their feed costs, since they also receive revenues from the sale of hides. The firm only needs to be concerned with whether or not the total revenues from both the beef and the hide markets are sufficient to recover the joint costs of cattle feed.

In attempting to analyze prices of one service relative to its costs, joint costs create considerable difficulty and controversy. One solution is to allocate a reasonable share of the joint costs to each of the joint products. Unfortunately, as economic theory demonstrates, there is no unequivocally correct way to allocate these costs among the various services. For example, the factors that determine the level of loop costs (geographic characteristics of the service territory, customer density patterns, zoning requirements, technological changes, etc.) cannot be closely linked to the various customer classes or service categories. As a result, there is no unambiguous "cost causative" method available for allocating the loop costs; the most to be expected is an allocation method that produces reasonable and equitable results.

Ideally, the allocation method selected to allocate loop costs would recognize the relative strength of demand for the joint products that make use of the loop, consistent with the method by which joint costs are recovered in competitive markets. For example, the toll category would bear a share of the loop costs which is greater than its proportion of usage, in order to reflect the greater value associated with longer distance calls, and the higher cost of substitutes for these calls relative to local calls. (The substitute for a local call may be driving across town to talk; the substitute for a toll call may be flying across the country to meet).

Differences in value have historically been reflected in the way many jurisdictions price

local and toll service, even though regulators many not have explicitly based their decisions on the underlying theory of joint cost recovery. In some jurisdictions, regulators have been urged to rely more heavily on fully allocated cost studies. The latter studies can derive widely varying results, and thus the direction in which prices will be moved by such an approach will depend heavily upon the particular allocation formula selected. For example, if loop costs are allocated on the basis of relative minutes of use, without weighting the toll minutes more heavily to reflect their greater value, the result will be to shift costs away from the toll market, placing more of the joint costs on the local market.

The concept of marginal cost

Marginal cost is one of the most important concepts in standard microeconomic theory. It focuses attention not on the total level of cost, nor the average level of cost (concepts that are often more familiar to non-economists) but rather on the change in costs that occurs as the volume of output is increased or decreased.
Marginal cost is defined as the change in the total cost of production resulting from an extremely small change (upward or downward) in the level of output. To be strictly technical about it, marginal cost is the first derivative of the total cost function with respect to output. It can be recognized at once that the minimal measurable change can be extremely small--one more milliwatt of electricity, one more drop of water, one more cubic foot of natural gas, one more second of calling duration, or one more local loop. And, when dealing with the pure theoretical concept (as measured by the first derivative) marginal cost is defined as the rate of change in total cost as volume changes by an even smaller amount--an infinitesimally small amount.

Practical problems with estimating marginal cost.

In attempting to estimate marginal costs, the analyst often encounters practical difficulties when the measurements are directly calculated at the smallest possible level. Accordingly, most practical estimates of marginal cost are based at least in part on a slightly larger increment of output than what is envisioned in economic theory.
A marginal cost study will typically measure the average rate of change in total cost across a moderate size incremental change in output. For instance, the analyst might determine what happens to total cost as output is varied in a range from 90% to 95% to 100% to 105% to 110% of current production. Then, the analyst will average or smooth the results, in order to estimate the rate of change in cost which is occurring within this limited range. While our theoretical objective is to estimate the cost of adding one more call, one more minute of use, or one more loop, in practice these are not large enough increments to provide reliable results, if calculated directly. Because various components of the technology can only be obtained in specific sizes or lumps (e.g. 50 or 100 pair cable is manufactured, but not 63 pair and 64 pair), if one directly checks the change in total cost as one more call is placed, or one more loop is added to the network, the result will generally be zero (and occasionally will be a very large number). In order to overcome this sort of "lumpiness" and for other practical reasons, the analysis must necessarily include some degree of smoothing, or averaging, of the cost data over a slightly larger volume of loops (such as the entire range from 50 to 100).
Accordingly, the marginal cost of a local loop will actually be figured as the additional cost of a modest increment of investment divided by the number of loops encompassed

by that investment (e.g., if 300 more loops can be added for $3,000, the cost per loop is $10). The precise size of the increment being studied is a matter of judgment, taking into account the lumpiness of the various components of the service. In turn, these characteristics depend upon the particular technology being used, and the manufacturing practices of the suppliers of that technology. For instance, assume that copper wires are manufactured in cables of 25, 50, 100, 300, 600, 900, 1200, and 2400 pairs. Under these circumstances, any attempt to measure the increase in cost associated with a change from 723 to 724 pairs will be fraught with difficulty. One solution is to focus on the entire increment from 600 to 900 pairs. For instance, one can take the difference in cost between the 600 pair cable and the 900 pair cable, and divide by the change in the number of pairs (300). The resulting cost per pair is the (approximate, or smoothed) marginal cost per pair within this range.

How marginal costs can properly be estimated in practice

One approach to marginal costing is to measure the change in the total cost of an individual cost component as its quantity is varied by small finite increments well beyond the zero range. The marginal cost of this particular component can later be added to the marginal cost of other components, in a "building block" fashion, in order to derive an estimate of the marginal cost of a particular service.

Alternatively, one can build a model of the cost function, whereby it become feasible to first estimate the total cost of providing an assumed quantity of service (A), then estimate the total cost of providing a slightly different quantity of service (B), and then divide the cost difference by the change in the number of units (C). That is, (B-A)/C. Due to lumpiness, it may be necessary to repeat this process for a series of small increments (e.g. 1% or 5% variations in output), then average or smooth the data, in order to produce a reliable estimate of marginal cost.

No matter how elaborate the cost models, or how complex the study, the final result should generally be consistent with this basic concept of marginal cost; otherwise, the results cannot properly be considered an estimate of marginal cost. To the extent the model greatly deviates from this basic concept, the results cannot properly be considered an indication of marginal cost. If the model fails to distinguish between small changes in the rate of output and extremely large increments of output, the results cannot properly be described as "marginal cost."

Using the feeder/distribution cable for local loop on a specified route as an example, one might estimate the cost of installing sufficient feeder/distribution cable to meet projected future demand. Then, one might calculate the cost of a larger amount of cable along the same route which is sufficient to handle 105% of the projected future demand. The difference in these two costs would then be divided by the 5% difference in the number of loops under the two scenarios, to derive the marginal cost associated with a slightly larger volume of output.

Other approaches would also be reasonable; for instance, one could again start with the projected future level of demand along a particular route, then vary this downward by estimating the cost of serving 95% of the projected demand. Presumably, this would result in a smaller amount of cable being required, and thus a smaller level of cost. The difference in cost divided by the difference in the number of loops would equal the estimated marginal cost.

However, if the analyst considers the change in total cost associated with adding or deleting the route in its entirety, the resulting cost figure will have little or no resemblance to marginal cost. While such a cost estimate could arguably be described as a form of "incremental cost" (where the increment in question is the particular route, or the customers located along this route), it will not approximate marginal cost. To the contrary, the computed figure is likely to be very similar to average total cost.

In spite of data smoothing or averaging of the data, is marginal cost estimate actually be the same as average cost?

Not if the analysis is handled properly. While some averaging is necessary, because of the necessity of working with data for a range of output greater than a single loop, it is vitally important not to move to the opposite extreme. If the study focuses on a very large increment of output, the results will not provide a reliable approximation of marginal cost. Most significantly, if the study is to provide an estimate of marginal cost, it must not encompass the change in costs as output increases from zero. If it does, the study will greatly deviate from the theoretical definition of marginal cost, and the results will be completely unreliable for any purpose where an estimate of marginal cost is desired. In fact, if the study analyzes the change in cost as volume increases from zero to the total amount of output, the study will tend to approximate average total cost, rather than marginal cost. This is what occurs in the example mentioned previously, where the entire cost of placing a feeder and distribution cable is considered, rather than the rate of change in this cost, as the size of the cable is varied.
For most purposes close approximations of marginal cost are far more useful than average total cost. Parenthetically, it should be noted that with relatively little additional effort a study designed to calculate marginal cost can be expanded to derive estimates of average cost, as well.

Fixed and sunk costs, as they relate to marginal and incremental costs.

Fixed costs are simply those elements of the firm's total cost which do not increase as the volume of output increases. The difference between fixed costs and sunk costs is that the former can be reduced or eliminated if the firm is willing to exit the market entirely (e.g., by converting its equipment over to another purpose). In contrast, sunk costs cannot be avoided or changed even by discontinuing production entirely; thus, they are considered irrelevant for most economic decisions. A simple example of a fixed cost is the cost of owning a factory building; as long as the building is in use as a factory, its costs are unavoidable (and they do not vary with the volume of output produced by the factory). However, if the firm discontinues production, and sells the building to someone who converts it to another use, it will avoid the costs of ownership. Hence, the cost is fixed, but it is not sunk because the building can be readily converted to another purpose.
A simple example of a sunk cost is the cost of writing a novel. Once this cost is incurred, it cannot be avoided, reduced, or eliminated, regardless of whether or not the novel is published, or how many copies are sold. Stated another way, sunk costs are irretrievable once the decision to incur them is implemented. From that time forward, they are completely irrelevant to any pricing, production, or other economic decisions that must be made.
In telecommunications, the cost of installing buried copper cable is a sunk cost: once the

cable is in place, no future decision will alter those installation costs, or allow them to be not incurred. If the company is able to salvage some of the material involved, the salvageable portion of the cable cost is considered fixed, but not sunk. However, the labor needed to engineer and install the facilities is irretrievable. Therefore, once the labor costs of installation have been incurred, they are irrelevant to future decisions about the appropriate price level for the service or services that utilize the cable.
In the calculation of marginal or incremental cost, fixed and sunk costs are canceled out in the computations. This is one of the most distinctive attributes of the economist's concept of marginal cost, setting this concept apart from more conventional notions of average or total cost. The reason for this distinctive treatment is straightforward: since fixed and sunk costs do not change with the volume of output, they have no direct impact on the level of marginal cost, which is the change in total cost associated with a change in output.
Economic theory suggests that marginal costs are of particular importance in establishing prices, and thus fixed and sunk costs are of little or no relevance to pricing decisions. More specifically, economic theory demonstrates that since sunk costs are unavoidable, they are irrelevant to pricing decisions. Fixed costs are also irrelevant, to the extent the firm intends to continue in operation, and thus cannot avoid the costs in question. For decisions concerning whether or not to continue in operation (or whether or not to enter a market, or to produce a particular product or service) fixed costs can be significant, and must be considered, if those decisions determine whether or not the fixed costs in question will be incurred. For instance, if an entire factory can be sold or rented to someone else if a particular product line is discontinued, then the fixed cost of the factory must be considered in evaluating whether or not to continue to sell the product in question. But, once the decision to incur fixed costs is made, these costs do not directly affect the level of marginal cost, and thus fixed costs are largely irrelevant to pricing decisions. Hence, if the firm decides to continue to produce the item in question, the fixed costs of the factory should not influence the decisions it makes concerning the optimal price of that product. Thus, for example, it should focus on marginal cost data for pricing purposes (because this excludes the fixed costs) rather than average total cost (which includes fixed costs).

The concept of the 'run.'

The "run" is crucially important in determining the level of marginal cost, because the "run" determines the extent to which costs are fixed or variable. Mathematically, the costs of fixed inputs are canceled in the calculation of marginal or incremental cost. Put another way, since marginal cost is the change in cost associated with a change in output, those elements of cost which remain fixed as output varies will have no direct impact on marginal cost.
Thus it is very important in any marginal cost study to know which inputs are considered fixed (or sunk) and which are considered variable. To a large extent this is a function of the particular planning horizon, or "run" which is selected by the analyst. In fact, the same item may be considered variable in the "long run" but fixed in the "short run."

The 'short run.'

In the short run, while some inputs are variable, most inputs, particularly the size and mix

of the firm's plant and equipment, are considered fixed. They do not vary with the level of output. These fixed costs continue even in the event of zero output, unless the stoppage is of a permanent nature, such as the liquidation of assets and concomitant obligations. In contrast, variable costs increase as output increases (although not necessarily proportionately) and decline as output is reduced. Among the costs generally considered variable are materials and labor, as well as any equipment that can quickly and easily be installed and removed or that is readily reusable for other purposes.
The short run thus corresponds fairly closely to the world in which a firm operates on a day-to-day basis. Managers realize that they have existing facilities. These facilities are taken as a given and only rarely is consideration given to the possibility of selling, abandoning, or dismantling these facilities. The variations in cost that are of particular interest are those corresponding to the routine operational modifications that the firm can readily implement in response to changes in demand, given the existing facilities.

The 'long run.'

The long run is a more abstract concept: a theoretical planning horizon where most inputs are variable, including the scale and type of plant used by the firm. A cost which is considered fixed in the short run may be regarded as variable under a long-run planning horizon. When considering long run production decisions, the firm can analyze virtually any size of plant and mix of inputs (e.g., copper vs. fiber)--a luxury not available in the short run. With this greater flexibility, it is often feasible to produce additional output at a lower total cost in the long run than is possible in the short run, where fewer options are open. For a firm that already has a heavy commitment to certain technologies, the long run may be somewhat less relevant than the short run. Conversely, to a competitor considering whether or not to enter a particular market, the long-run planning horizon is more relevant than the short run.
Some people make the mistake of basing the distinction between long run and short run on calendar time. While there is some correlation between the passage of time and the degree to which an input is fixed or variable, the concepts are not identical and should not be confused. Rather, it is the extent to which the firm has flexibility in selecting and operating its plant and equipment which really determines this distinction. For a firm which has not yet entered the industry, the long run planning horizon, and associated cost curves, are of primary relevance. In contrast, for many decisions concerning a firm that is already operating in the industry, a short run planning horizon is likely to be of particular interest, since this will reflect various past investment decisions, which can be taken as given.

Incremental cost, and how this concept differs from marginal cost.

While marginal cost has a very precise meaning in economic theory, incremental cost is a much broader concept, which encompasses a wide variety of different cost calculations. At one end of the spectrum, incremental cost can be viewed as an "average" level of marginal cost, if it is computed over a narrow increment in the immediate vicinity of the current volume of production. For instance, if one calculates the total cost of producing the existing volume of output, then calculates the total cost of producing 105% of this volume, then subtracts the former amount from the latter and divides by the 5% change in the number of units, the result will be the incremental cost per unit associated with a 5%

increase in volume. Depending upon the nature of the cost function, this incremental cost may be a very close approximation of the "average" level of marginal cost which occurs over this narrow range (from 100% to 105% of the existing production level).
At the other end of the spectrum, if the increment encompasses the entire range from zero to the total volume being produced, then the incremental cost will be equal to average total cost (and it will not approximate marginal cost).

Specific terms that have been developed, which describe various types of Incremental Cost.

In regulatory proceedings in other jurisdictions where telecommunications cost methodologies have been discussed, it became apparent that it was confusing to use the same term, "incremental cost," to describe a wide variety of different cost concepts. Accordingly, an effort was made by various experts to distinguish the various different concepts, and agree upon specific terminology to describe each concept.
This effort to reach a consensus concerning appropriate terminology was reflected in an opinion and order issued by the Colorado Public Service Commission after a comprehensive investigation of costing and pricing issues as they related to the telephone industry, in Docket No. 92R-596T. This order outlines four distinct definitions of incremental cost. While some of these distinctions are fairly subtle, without this sort of clear terminology and careful thinking, a lot of the legitimate controversy is lost in a sea of confusion. The most general definition reads as follows:
Total Incremental Cost. The change in total cost resulting from an increase or decrease in output. In mathematical terms, total incremental cost equals total cost assuming the increment is produced, minus total cost assuming the increment is not produced. [Decision No. C93-443, Public Utilities Commission of Colorado, Docket No. 92R-596T, p.7.]
The type of incremental cost which was preferred by some of the telephone industry witnesses in the Colorado proceeding is called Total Service Long Run Incremental Cost (TSLRIC), and is defined as follows:
Total Service Long Run Incremental Cost. Total service long run incremental cost is equal to the firm's total cost of producing all of its services assuming the service (or group of services) in question is offered minus the firm's total cost of producing all of its services excluding the service (or group of services) in question. [Ibid.]
The Colorado order includes an extended discussion of the TSLRIC method, including the following explanations:
Total service long run incremental cost includes both fixed and variable costs specific to the service (or group of services) in question.
The total service long run incremental cost for a group of services is at least equal to the sum of the total service long run incremental costs of the individual services within the group. If the total service long run incremental cost for the group is greater that this sum, the difference is equal to the shared costs attributable to the group of services. In other words, these, shared costs are part of the total service long run incremental cost to the group but are not part of the total service long run incremental cost of any individual service within the group. [Ibid. p. 7-8.]
The Colorado decision includes another definition which essentially, restates the TSLRIC on a per-unit basis (e.g., per loop or per minute):

Average Service Long-run Incremental Cost. The total service long-run incremental cost divided by the total number of units of the service. [Ibid. p. 2.]
Neither the TSLRIC nor the ASLRIC approximates marginal cost. That theoretical cost concept is most closely approximated by the Incremental Service Incremental Cost (ISIC), which is described in the Colorado decision as follows:
Incremental Service Incremental Cost. The change in total cost resulting form increasing (or decreasing) the quantity of output of a service by a small number of units, divided by that small number. If the cost function is smooth and the increment is sufficiently small, incremental service cost will approximate marginal cost. [Ibid. p. 4.]

Total service incremental cost and the related concept of average service incremental cost.

The "total service incremental cost" has been advanced by various telephone industry costing experts as the most appropriate type of incremental cost for use in evaluating price levels for certain services. [Richard D. Emmerson, "Theoretical Foundation of Network Costs," INDETEC, Marginal Cost Techniques for Telephone Services: Symposium Proceedings, p. 149.] Total service incremental cost has been defined as "the change in total cost resulting from adding the entire amount of service or output to the company's total output with the levels of output for all other services remaining constant." [Alfred Kahn, The Economics of Regulation: Principles and Institutions, (New York, John Wiley and Sons, Inc., 1970, Volume 1.]
If the purpose is not to establish specific prices or test for subsidies, but rather to determine whether or not an entire service should be offered (added or discontinued), then the "Total Service Incremental Cost" (TSIC) may be useful. Thus, for example, if the firm (or regulators) are trying to determine whether or not it is profitable for the firm to offer voice mail service, the TSIC of voice mail service would be compared with the corresponding total incremental revenues that would be received from this service; if the costs exceed the revenues, entry into voice mail business would not appear to be appropriate. In contrast, if one is trying to decide what price to charge for a specific tariff item, and there is no question about whether or not the service will be offered, then the TSIC would not be especially useful. As a general principal, where pricing is the focus of interest, the more appropriate measure of cost is likely to be the "incremental service incremental cost," because this measure can provide a close approximation to marginal cost, if it is properly estimated.

Some of the problems which are often associated with incremental or marginal cost studies.

The most serious problem is that many of these studies do not come close to estimating marginal costs (or the incremental service incremental cost). In fact, studies offered by telephone utilities sometimes fail to even estimate any other precise, readily identifiable version of incremental cost, such as the average service incremental cost. To the contrary, despite being labeled "incremental cost," these studies often contain a hybrid mixture of different costing concepts.
For instance, some cost studies focus at least in part on current reproduction costs (the cost today of reproducing the existing mix of facilities), despite being labeled as "forward looking." In that case, they may be greatly overstated, since they reflect all the increases

resulting from inflation without reflecting the offsetting decreases resulting from technological improvements. Sometimes, studies provided by telephone utilities include an amalgamation of embedded, reproduction, and forward-looking costs, blended together and computed on an average cost basis--despite being labeled "incremental cost."

To the extent such studies are focused on incremental costs at all, the increment is usually so large that it encompasses everything from quantity zero through the entire existing output (or beyond). Hence, such studies actually measure average costs, not marginal cost. The essence of marginalism is that it focuses on the cutting edge, or margin of decision-making, rather than the total cost of production. For an incremental cost study to provide a useful approximation of marginal costs, the increment studied must be fairly small, and (most importantly) must not encompass output level zero. To be sure, if one is interested in knowing whether the total increase to revenues resulting from a new service will exceed the total increase to cost, some version of total service incremental cost may be of interest. However, if one is interested in pricing decisions as they relate to economic efficiency or profit maximization, the total service is too large an increment to be useful. Rather, one needs to look at a much smaller volume of output--coming as close to the pure theoretical definition of marginal cost as is practical.

Likewise, marginal cost information is useful in determining the presence or absence of subsidies. While the term "subsidy" is often used loosely to describe any situation in which a service appears to be priced below cost, under the economic definition a service is said to be subsidized only if its price is below marginal cost. When speaking of whether or not a particular item is subsidized (e.g., local service purchased by residential customers who could not afford to pay a higher price, and thus would otherwise not be on the system), the "Incremental Service Incremental Cost" is the relevant test for a subsidy. However, when analyzing whether or not an entire category of service is being subsidized in totality (e.g., basic local service as a whole), the Total Service Incremental Cost is generally the appropriate test for a subsidy.

The relevance of TSLRIC studies for purposes of 'rate rebalancing' and universal service rate investigations, where claims of uneconomic subsidies have been made

TSLRIC studies are entirely appropriate for this purpose, as opposed to fully allocated cost studies, which have been increasingly discredited by most sectors of the industry and most outside observers because their methodology is limited to embedded costs and fails to provide for an adequate depiction of future economic costs of telecommunications networks.

The proper treatment of joint costs in a TSLRIC study.

In a pure TSLRIC or TSIC approach, joint costs would be excluded, since those costs are necessary for the production of the other services and would still be incurred in the total absence of the service in question. However, since this claim can equally well be made for every other service offered by the firm, it is clear that the application of TSIC studies to issues of cost recovery and pricing will ultimately entail some sort of allocation process, similar to the treatment of costs in a fully allocated cost study, or else it will be necessary to provide markups above TSIC in developing prices that are sufficient to ensure recovery of the firm's joint costs. Similarly, if the analyst excludes common costs

from the TSIC study, it should be understood that recovery of these costs will require application of a markup above TSIC for pricing purposes, or it will be necessary to add an allocated share of common costs to the pure TSIC results.