DEVELOPING PRICE INDEXES FOR CONSULTING ENGINEERS: THE CANADIAN EXPERIENCE

VOORBURG GROUP CONFERENCE
SYDNEY, AUSTRALIA
OCTOBER 17-21, 1994

STATISTICS CANADA

September 1994

DEVELOPING PRICE INDEXES FOR CONSULTLING ENGINEERS: THE CANADIAN EXPERIENCE

In 1990, Statistics Canada began development of price indexes for the output of consulting engineers, as part of its programme of improving statistics in the service sector. Although none of these problems is unique to service industries, measuring the price change of services is generally considered more difficult than for goods, as outputs of many service industries are not so easy to observe and define, and often these outputs are not reproduced exactly from one period to another. Furthermore, the statistics on outputs that enable samples to be drawn for particular commodities and provide aggregation structures to combine the price measures for those commodities may not be so well developed for services as they are for goods.

All these problems exist to some extent in pricing engineers' work. A paper presented to the Voorburg Conference of 1991 "The pricing of services of consulting engineers", described what had been done in defining outputs, and the approach that would be used for getting price comparisons over time. By this method, "model pricing", which had been used successfully in construction and in some goods areas, each selected producer would be asked to estimate what price they would expect in the current market for a piece of work, a particular project, that had been carried out in the past. By comparing this expected price with the actual price obtained in the past a measure of price change over that period is obtained. Indexes would be calculated by combining the different measures from different respondents. It would be necessary that the individual model projects chosen should reflect the diversity of activities to the extent that the diversity affected the rate of price change. At first we did not know which influences would be the most important on price change. The survey of fees earned in the industry distinguished among revenues for different services provided, from different fields of specialisation, and from different regions of Canada (see appendix 1 for details). A large part of the paper was concerned with how representative the sample chosen at that time was in reflecting each of those revenue patterns.

A year later, when respondents had begun to supply their estimates of price change, it was becoming clear that the model pricing approach would need to be modified. The concerns were mentioned in a brief update to the 1992 Voorburg Conference. This paper reports on the modifications made to the model pricing method, and, with the publication of the first sets of indexes in the spring of 1994, explains the other statistical decisions made in the design of this survey that may be of some value to others in the same field.

Model pricing and its modification

A problem for measuring prices in the consulting engineering industry is that each project which occurs is unique. No one provides the same services twice, as the requirements of each contract differ. This problem has occurred in measuring price change for construction work, and in some areas of heavy equipment manufacturing. Price index makers have dealt with it by pretending that the work could be reproduced, and getting the producers to estimate what the price would

be. This has been successful in construction and manufacturing.

In construction, model pricing works this way. Any construction project is the result of a number of pieces of work by specialised trades, excavation, carpentry, plumbing, electrical, etc., whose combination is looked after by a general contractor. When a general contractor is bidding on any project he includes estimates of what each individual trade will be costing, based on sub-contractors' prices for materials, labour and their profits, and adds to it his costs and an expectation of profit. Since contractors in each of these trades bid with the general contractor on any project, it is fairly straightforward to get them to provide hypothetical estimates to a survey taker on the same unchanging job specification from time to time. The influences on prices are the costs of materials and labour, the way in which they are combined to produce various products, and the current market conditions which will affect the expected profit. Although in construction the exact same product is not usually made twice, the work that is done uses the same inputs. Any technological improvements apply equally to most processes, and companies can move easily from one buyer to another. It is possible therefore to combine the reported price movements for individual trades in different ways to model the overall price movements for different projects.

In engineering there are some differences. Although a project usually combines several different pieces of work, all the work is more dependent on the skills of the firm's employees. There is not the same activity of converting materials into something else. Consequently, when a company is costing out a project, it proceeds by estimating the number of hours of work that will be required by each of their specialist staff assigned to the various tasks that make up the project. The sum of these hours is the "total billable hours". When the salaries of the various staff are applied to the hours, they get an estimate of direct labour cost. Companies then multiply this figure by a certain amount to get their estimate for the value of the project. They usually have a sense of what this multiplier should normally be - it is enough to cover all the other costs they incur which cannot be charged directly to the project, but which have to be covered by the revenue - ancillary direct labour costs, the costs of other staff providing general services, overheads, depreciation on capital equipment, and a margin for profit. However, they will tend to adjust the multiplier as market conditions change, and as their cost situation changes.

In applying model pricing to consulting engineering we selected representative projects with each respondent and asked them to reproduce the way these projects were originally estimated. This involved getting records on the numbers of hours billed for various kinds of labour, and the amounts by which the total direct labour costs were marked up. In the re-estimating, we expected to find changes in three areas: the number of billable hours required, the salaries paid, and the multipliers by which the direct labour costs were marked up. We also expected that, although the multiplier covered the project as a whole, other costs, since they were separately identified in the detail of the work, might show different rates of change. In that case we would be able to develop different indexes for the different services provided.

In fact, the first re-pricing brought results different from these expectations.

Although labour rates might move slightly differently between different groups the differences were not enough to affect the estimations. We saw a lot of detail which all showed the same percentage change. There was no case of the numbers of hours required being changed, nor of the mix of hours of different grades changing. Most seriously, the multipliers did not show as much change as we would have expected. To deal with these problems and to make best use of the records that respondents normally keep (some respondents, while cooperating in the survey, felt it included a lot of unnecessary paperwork), the method was modified to collect observations on the different influences independently. The indexes now are calculated by combining independent measures of changes in wage rates, in multipliers, and in direct labour requirements.

The first repricing also taught us something about the relative importance of different influences on price change. To most companies, the particular service being provided was not so important. They could switch easily from doing one thing, say design, to doing another, advisory service, perhaps. On the other hand, partly because of the expertise they had acquired in particular areas, and partly because of relationships built up with particular clients, it was not so easy to switch whole projects from one area to another. It appeared, therefore, that price indexes would not differ much by the type of service provided so much as they would in different parts of the country, and particularly, in different fields of specialisation. The design of the survey proceeded on that basis.

There are now three connected surveys which measure the three influences on price change independently; one on labour rates, which applies company by company, one on multipliers which applies by company in specific markets, and one on the change in labour requirements, which at the moment applies to the industry as a whole.

For labour rates we ask what is the change in the rate paid for someone with the same experience doing the same work, compared to the previous year. Most companies review their salary scales annually, so have little problem in supplying this information; they normally provide it at the beginning of the coming year. If the salary norms are reviewed during the year, or if additional bonuses are paid across the board, the rates would be revised. In the climate of the early nineties this did not happen.

In observing markups we moved from estimates of what companies expected to their records of what had actually happened. We found that in the model pricing exercise the same percentage markups would be quoted even though the market situation had clearly gotten worse. The explanation was that although the expected markup from earlier might not have been achieved, it was still the basis on which they would expect to contract for future work. In this industry, in Canada, amendments to the terms of payment would be made after the work had begun. If the client was having trouble, the engineering consultant would take less to keep the project going. In effect model pricing was giving us book values rather than reflecting what would probably be earned.

Almost all companies keep records of their ratios of revenue to salary costs. It helps them monitor their performance. The ratios different companies keep vary

slightly in their definition and are called by various names, but they all track how their revenue from projects relates to the salary costs billed to those projects. The term "realised net multiplier" is used here. It is similar to the ratio which in model pricing we were asking respondents to estimate, under a set of fixed hypothetical circumstances. The main difference in concept is that it records after the fact the markup actually achieved, whereas in the model pricing we sought, unsuccessfully, what would be expected to be achieved. Except for the delay in getting data, this is an improvement. The practical drawback of using their actual records is that they cover all projects, so a decline, for example, in the realised net multiplier might occur not because the multipliers have fallen because of anything connected with prices, but because in the second period there were more projects which typically had lower multipliers than there had been in the first period - the standard defect of all measures of average unit values. We tried to mitigate this drawback by eliminating as much of the changing mix as possible.

It had been discovered that price variation was most dependent on the market - which meant region of the country, and/or the kind of industry purchasing the work. For the industrial specialisations, mining, pulp and paper, oil and gas, power generation and transmission, and other industrial services, the market appeared to be national. For other engineering work, buildings, transportation facilities, municipal services and environmental services, the markets were geographically smaller, so separate markets were defined for individual provinces, or groups of provinces. Respondents were asked to provide realised net multipliers for each of these market categories in which they were active (see Appendix 2 for the list of markets). Most companies reported on the three or four that were their main sources of revenue. In some cases where they claimed that two or more separately identified markets were the same so far as their pricing experience was concerned, the combined ratios were accepted as observations for each of those markets.

In this way, the variation in markups due to operating in different markets was eliminated, and the groups of projects were more homogeneous. In addition, if some unusual work was carried out, the ratios have been reworked excluding that abnormal contract. To do this effectively requires close contacts between the survey officers and respondents, but that has been easier to achieve with the greater confidence respondents now have in the methodology.

To measure labour productivity change the model pricing principle has been kept. Some aspects of productivity change, the more efficient use of machinery, reduction in overheads, and so on, will be reflected in changes in the realised net multipliers. There is no a priori way to know how the multipliers will tend to move over time. A lowering of overhead costs will tend to lower the multiplier. But an increase in fixed capital will tend to raise it. Even if the acquisition of more capital leaves the amount of labour required unchanged, depreciation will rise. But if, in addition, the extra capital means that less hours of labour input is required, in the absence of salary rate changes, the ratio of revenue to labour costs will rise even further, as the denominator in the fraction is lower. To capture the missing influence we need to measure the change in labour requirements over time. To do this we go back to the original sample projects and ask those respondents willing to give the time to it to re-estimate the quantities required. This is being done on a

rotating five year cycle, each respondent is asked to re-cost a project every five years. It is difficult to observe small incremental changes every year, the longer period between reviews makes them easier to see. However, it also makes it more likely that we are told that such and such a project would never be done today. Perhaps extra environmental tests are required, or more design work would be necessary. In such circumstances respondents have the choice of estimating the extra work that would be required and comparing those costs. Consistent with the view of the industry that it is flexible in the products it supplies, the productivity estimates are assumed to be the same for all markets.

Questions of index estimation

In this paper so far the discussion has been about obtaining as accurate information as possible on the change in price from one respondent, in one set of circumstances. To produce indexes the individual observations have to be combined in some fashion. It is useful to distinguish in this process two levels of index number construction. We have defined a set of indexes for each of the markets identified in appendix 2. With each of these indexes is an associated weight, so that aggregations above this level can be done according to textbook theory. Below this level, however, practical problems intrude on the theory. It is not easy to find the appropriate weights, and data are missing; samples do not stay the same from one period to another and adjustments must be made for these discontinuities.

Estimation at the detailed level

The index for each defined market is estimated by multiplying together the three subindexes for labour rates, realised net multipliers, and labour productivity. The labour productivity index is the same for each market, but the other two are each calculated from the observations that apply to that market. The set of observations may not be the same for each of those sub-indexes, as labour rates are available ahead of time, but the multipliers are only available at the end of each company's accounting year. As different companies have different accounting years some observations will be missing for the first calculation.

At this level of detail there is a problem in deciding how to combine the observations from different companies. Because all companies come and go fairly swiftly in any given market, it is not possible to compare current observations to base period values. In fact, comparisons have to be made over short periods of time, chaining the movements from year to year. Arithmetic means of period to period price relatives are not transitive. If one price only is reduced in period two, and returns to its original level in period three, the index in period three will not be the same as in period one, but will be a little higher. This is because the weights are implicitly changed to give more weight to the reduced price when it increases again in period three. Consequently geometric means of the individual price relatives were used to calculate the indexes.

The choice of weights was more difficult. Larger companies account for more volume than smaller ones, though we do not know whether they show systematically different price behaviour. The ratio of the size of the largest

respondents to the smaller ones is even greater than the relative shares of large and small companies, as we do not know the values in any given market of the unsampled smaller companies. Also, it is quite likely, for the realised net multipliers, that the observations from small companies are more reliable indicators of overall price change than those from larger ones. As the larger companies have a broader range of projects there is more chance that the average change in the multiplier may be affected by a shift in the mix of projects from one year to the next. Finally, as companies move in and out of markets from one year to the next, it is not easy to establish firm weights. There is a danger of overweighting companies in the following year and introducing the same kind of bias into the indexes as using arithmetic averages does. For all of these reasons, it was decided to use unweighted averages to combine individual price movements at the most detailed level. (With one exception: when a company provided one set of data covering n markets together, its observation in each market was given a weight of 1/n.)

Review of results

Prior to the development of these indexes, estimates of price change in this industry were based on the movement of earnings, specifically of average weekly earnings in a group of industries which included consulting engineers as well as architects and technical services. There should be two benefits from developing direct price measures rather than relying on the measures based on earnings figures. The first is that the price indexes should reflect market changes more closely. The second is that the long-term trend should be more accurate; in this case that the output price indexes should move less than the earnings figures which cannot reflect the savings from productivity increases over time.

The first indexes published cover the period 1989 to 1993 during which the economy suffered a major recession. The variation of index movements in different markets certainly suggests that differences in short-term market conditions were captured well. Generally, the impact of the recession was felt in the private sector - industrial applications and buildings, before the public sector - largely transportation and other civil engineering. It was felt in different degrees in different parts of the country, the depressive impact on prices being greater in eastern Canada than in the west, and prices being slower to recover in the east in 1993. But although these different measures are interesting to the industry and their clients in helping them monitor individual market conditions, for the main purpose of these indexes, measurement of the overall price change in the industry, the appropriate comparison is between the total index and the behaviour of earnings figures.

Comparison of earnings figures and output price indexes 1989 to 1993

	Average weekly earnings		Output price index	
	Index	% change from previous year	Index	% change from previous year
1989	82.7		92.5	
1990	90.1	9.0	95.1	2.8
1991	95.9	6.4	98.2	3.3
1992	100.0	4.3	100.0	1.8
1993	102.5	2.5	100.3	0.3

Clearly the price indexes reflected the impact of the recession better than the earnings figures. We know that average earnings data do not track the impact on markets well during economic declines as there is a tendency for lower paid employees to be laid off first. A paper by Pierre Gagnon examined this in the engineering industry from 1989 to 1992. He showed that when the 1989 mix of engineers of different grades was applied to earnings figures for 1990, the increase in average earnings was higher than that shown by the changed mix, but in 1991 and particularly in 1992, when employment was falling, the fixed weighted average was lower. Nevertheless, even allowing for that effect, the price indexes showed lower changes throughout the period of recession, and into the early stages of recovery.

Whether these indexes will show an overall long term lower trend has yet to be seen. We do not yet have a comparison of the two measurements during period of expansion. Just as earnings figures underestimate price slowdowns during recessions, they may well underestimate increases during expansions, so we could expect to see the price indexes moving up faster than average earnings in the future. At present the difference in accumulated price movement, 8% as against 24% since 1989, is substantial, and should not be expected to be maintained. A comparison between input and output indexes for construction shows a difference of about 1% a year overall over the last twenty years and about 2% between the output indexes and wage rate movements.

Summary and conclusions

The survey has been successful to the extent that it is now in regular production, and has produced figures that are credible to the industry and to our main clients in the System of National Accounts. It is not being produced in the way we had ideally planned at the beginning. In retrospect we lost some time in trying to apply a particular approach to price measurement before we knew enough about the organization of the industry, how prices are set and changed, and about the records that the industry keeps. However, even had we known, it would have been necessary to have collected prices as a pilot survey for a while to see what practically were the main influences on price change. The need to distinguish price movements for different markets rather than for different products might have

been guessed from considering the process by which this industry creates products, but it had to be established empirically.

The decisions made to deal with some estimation problems were taken in the context of what information was available in this Canadian situation. In different contexts the solutions may have been different. However, the selection of unweighted geometric means at the most detailed level of estimation was the first application of this choice. Geometric means at the most detailed level have much to recommend them in general situations (see Bohdan Schultze's paper on the subject). It will likely be the norm which would only be rejected for special reasons in similar future surveys.

References:

B. Schultze, "Price Indexes below the basic aggregation level", Bulletin of Labour Statistics, vol 2, 1987, IX-XVI

P. Gagnon, "Prices, Earnings and Employment in Engineering Services", working paper, Business and Trade Statistics analytical conference, Statistics Canada, 1993

R. Lowe, "The Pricing of Services of consulting Engineers", working paper, Voorburg conference 1991, and update 1992, Statistics Canada.

Appendix 1: Categories of fee revenue.

By type of service:

Advisory services- environmental or other
Design services- environmental or other
Construction management
Project management services
Other

By field of specialisation:

Buildings (structural)
Buildings (mechanical and electrical)
Highways, bridges, tunnels, railways
Transportation facilities
Municipal (roads, streets, water supply)
Municipal (sewage treatment, waste disposal)
Other environmental services
Mining, metallurgy and primary metals
Pulp and paper
Oil, petroleum and natural gas
Power generation and transmission
Other industrial
Other

By region:

Province, or if foreign, by country

Appendix 2: Most detailed areas for which individual indexes are designed (Canada only)

	Atlantic provinces	Quebec	Ontario	Manitoba/ Saskatchewan	Alberta	British Columbia	Canada
Buildings (str.)	x	x	x	x	x	x	
Buildings (m+e)	x	x	x	x	x	x	
Highways etc.	x	x	x	x	x	x	
Trans. facilities	x	x	x	x	x	x	
Municipal (roads)	x	x	x	x	x	x	
Municipal (waste)	x	x	x	x	x	x	
Environmental	x	x	x	x	x	x	
Mining							x
Pulp and paper							x
Oil and gas							x
Power gen.							x
Other industrial							x
Other services							x