

Paradigm Lost

J. E. Everett

*Centre for Exploration Targeting,
School of Earth and Environment
University of Western Australia, Perth, Australia*

Jim.Everett@uwa.edu.au

Abstract

As an academic consultant advising industry, two contrasting experiences are frequent. One is that the industry practitioners have strong domain knowledge that the consultant lacks. The other is that the industry practitioners often suffer from paradigm paralysis, using models of thinking that are either fundamentally inappropriate, or, perhaps more commonly, have become inappropriate through the passage of time and changes in technology and computing power. The data available has often grown considerably. Despite this increase in data (or perhaps because of it) paradigm paralysis can cause decision makers to fail to make full use of the available information or, worse, can lead to outcomes contrary to those intended.

By coming in as an outsider, academic consultants can provide benefit by introducing new paradigms or updating previous paradigms, subject to the important proviso that they become familiar with the domain knowledge and work closely with the industry practitioners.

In this paper, the author discusses a number of projects, in which he has been involved, that illustrate the occurrence of paradigm paralysis.

Keywords: Decision Support, Paradigm, Model, Consulting, Industry, Mining.

Introduction

There is a story, perhaps apocryphal, that until after the Second World War, whenever a field gun was fired in the British Army, a man had to stand so many yards to the right. No one questioned this instruction, which was featured in each edition of the Standing Orders, until someone went to the Army Library and looked through all the previous editions. It turned out that, until the Boer War (at the end of the nineteenth century), the man had held the horses so that they would not bolt when the gun was fired. The horses had gone, but the man had not.

Despite a thorough search, I have not been able to find a valid reference for this story, but it very well illustrates a situation that I have frequently come across.

Material published as part of this publication, either on-line or in print, is copyrighted by the Informing Science Institute. Permission to make digital or paper copy of part or all of these works for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage AND that copies 1) bear this notice in full and 2) give the full citation on the first page. It is permissible to abstract these works so long as credit is given. To copy in all other cases or to republish or to post on a server or to redistribute to lists requires specific permission and payment of a fee. Contact Publisher@InformingScience.org to request redistribution permission.

A “paradigm” is broadly a philosophical or theoretical framework. Our actions and decisions are guided by the information available to us, but this guidance is strongly influenced by our underlying paradigms, which may be explicit (such as a scientific theory) or implicit (such as the skills and habits acquired by a trained operator). Without paradigms, we could not convert information into

decision and action, but inappropriate or obsolete paradigms may lead us grievously astray.

Kuhn (1996) discussed how paradigms shift after a lengthy period of stasis. More recently, Gill (2010) considers the extreme reluctance to renounce existing beliefs, referring to it as “The Law of Abandoned Expertise.” As Gill states it, “We are extremely reluctant to accept new information that forces us to abandon any expertise we have already acquired.” The phenomenon of paradigm paralysis, defined by Wikipedia as “the inability or refusal to see beyond the current models of thinking” is not confined to revision of major scientific theories, such as the recognition of continental drift (Bullard, Everett, & Smith, 1965). The phenomenon also applies in everyday personal and organisational behavior, and this paper will discuss some examples of paradigm paralysis leading to inappropriate or inefficient extraction of information from available data. As the examples will show, paradigm paralysis is often sustained or perpetuated by confirmation bias: the “tendency for people to favor information that confirms their preconceptions or hypotheses regardless of whether the information is true” (Plous, 1993, p. 233). There are of course instances where a slightly less accurate but simpler paradigm may be retained; for example, for everyday situations we all use Newton’s laws instead of resorting to the slightly more accurate relativistic formulation. But the evolution and general availability of the computer has rendered obsolete many paradigms that are still in common use, especially paradigms that relate to information analysis and decision support.

This paper will discuss a few such examples that I have encountered. Each example considers a real situation where an accepted paradigm either impeded ongoing performance or caused an unfortunate outcome. In each case, the accepted paradigm needed replacing or updating. This could be because the underlying model on which the paradigm was based was initially wrong or incomplete, but more commonly because the world had changed without the paradigm being updated. In each example, the problematic paradigm being used will be identified and the informing process used to attempt to rectify it will be discussed. Each discussion includes an attempt to identify the factors that tended to help or hinder the necessary paradigm shifts. Although the anecdotal nature of the evidence presented means that the conclusions are tentative and situational, it is suggested that they may provide an insight into a generally applicable approach to paradigm shift. To this end, a tentative taxonomy will be suggested, classifying the examples according to the nature of the paradigm problem and the appropriate solutions.

Background

My history is that, after post-doctoral research in geophysics, I worked in the oil and mining industries for eight years. Because I had computer experience, in the days before such experience was common, I was given the task of running project evaluations, for example Monte Carlo simulations of the probability distribution for the profitability of an oil or gas prospect prior to drilling. This led me into arguing with accountants and losing arguments with accountants, so I completed a couple of degrees in economics and commerce at the local university (and found the accountants were generally right). The university was just starting an MBA course and wanted staff with industry experience to teach on it. So I re-joined academia to teach management science and information systems in a business school for thirty years. This combination of commercial and academic backgrounds gave me many fruitful opportunities to be a consultant to industry. The consulting opportunities generally related to the introduction, development, or revision of the client’s information and decision systems, with the opportunities that became available as computing power became cheaper, simpler, and more generally available.

In many ways, I found consulting to be an even more rewarding experience than teaching, over and above its ability to shrink the gap between academic and commercial pay. For a start, it is much easier to connect with one or a few folk who have called you in to help, rather than a large class of students of varying motivation, some of whom just need the credit points and ask whether

the material is examinable, or cannot see the relevance to some unknown future life. Some MBA students were the worst in this respect: even to the extent of saying “I don’t need to learn this. I am going to be a manager, so will be able to delegate someone to do it for me.”

So the majority of consulting experiences came as a welcome diversion, dealing with a real problem, jointly learning with people doing interesting things. Generally, I found they had specific domain knowledge that led to as much learning on my part as on theirs. Further, the work commonly provided real problems and real data, leading to publishable material, either in collaboration with the practitioners or in disguised and generalized form. This kind of applied research I found much preferable to the alternative of wasting time writing grant applications and then seeking un-motivated industry collaborators.

As an academic consultant, I found a number of factors contributing to success (or disaster) and a few ways of sniffing out the opportunities that enabled an outsider to contribute. I shall first summarize my experience of the success/disaster factors, and then spend the rest of the paper discussing opportunities that I have encountered in recognizing paradigm paralysis and confirmation bias.

Factors for Success or Failure of a Consulting Project

For an outsider called in to help on a problem, the first essential is to establish the identity of a project champion and mentor within the organisation. The champion has to have enough influence to make sure the work is taken seriously and enough time to keep abreast of the project as it develops. One of my most memorable disasters was a job for a small sheet metal working factory. The owner was very enthusiastic, but had to go overseas for two months. He passed the project on to two of his staff who were unaware of, or did not share, his objectives. On his return, he quite rightly was angry at the waste of time and effort.

Although the consultant has been called in to give advice and help solve a problem, the existing staff nearly always have considerably greater domain knowledge, about both the relevant technology and the social/cultural environment in which they are working. It is essential to respect this domain knowledge and to learn as much of it as possible as quickly as possible.

This leads to a third essential requirement, that the consultant work with the client rather than for the client. A recipe for disaster is to come in, do an appraisal, go away and work out a solution, and then deliver it. Working with, involving, and co-opting clients has multiple benefits, both technical and psychological. Their domain knowledge can help prevent silly and embarrassing outcomes; involved staff have an investment in making a solution work instead of reacting with the “not invented here” syndrome, and they become champions spreading acceptance of the work among their colleagues. Commonly they come up with suggestions for further system improvements. If such staff can be encouraged to think that the ideas for improvement were theirs, and the consultant has been merely a facilitator, then they have an investment in its successful implementation. Which brings me to another recommendation for success: if asked to solve a large problem it is wise to identify a smaller subsidiary problem and work on that first. Expansion of the project is then in response to the client’s pull rather than the consultant’s push. Often a certain reluctance is advisable, along the lines of “Do you really need that,” followed by “Oh well, perhaps we could add it if you all think it would help.”

Across a variety of industries, it is surprising how the opportunities and problems of one industry reappear in another, often in a not immediately obvious way. Transferring the solution across such environments can provide a helpful antidote to paradigm paralysis. The “Hospital Laundry Costing” and “Ore Input Selection” examples that follow will be seen to illustrate this fortunate opportunity of cross-fertilisation.

Examples

Cardboard Box Cutting

One of the earliest consulting projects I was called in to was for a small company manufacturing cardboard boxes (Everett, Gordon, & Jarvis, 1976). The request was to advise on the size and quantity of the rolls of paper that they should stock, to balance the cost of running out of paper required for their varying daily orders against the cost of tying up capital in excessively large stocks of paper. As for many such assignments, the problem is like a set of stacked Russian dolls: the necessary range of paper sizes depended upon the range of box sizes they were manufacturing. Revising the range of box dimensions could reduce the range of paper sizes needed to stock, but persuading customers to accept standardised sizes generated further problems. For example, they were supplying two different fish exporters with boxes of the same volume, but different dimensions. Standardising the box size would help, but that turned out to be difficult because the two customers used different sizes of pallet on which the boxes were stored and moved.

As is so often the case, it helped to find a smaller initial project of value. Each morning, the foreman decided what papers should be used for the daily production of cardboard box blanks. The papers rolled through a cutter, with rectangles being cut, either a single width or multiple, depending on the relative widths of the required rectangles and of the paper roll, as shown in Figure 1. Usually, the available paper was too wide, so there would be some wastage at the edge. For years, the scheduler had been choosing assignments so that the strip of paper wasted was as narrow as possible: for example, for cuts 200 mm wide he would assign a paper width of 250 mm (wasting 50 mm on a single cut), instead of 700 mm (wasting 100 mm on a triple cut), even though the proportion of paper wasted on the triple cut would have been less (14.3% instead of 20%) and the equipment usage with the multiple cut would have been more efficient.




Figure 1: Cutting the blanks for cardboard boxes

Once the problem had been identified, the technical solution was obvious, but implementing it without embarrassing the foreman was more problematic. Fortunately the internal champion came up with a bright idea: he persuaded the foreman that the daily scheduling task was too routine for his skills and assigned it to an apprentice to whom he explained the revised procedure.

Railway Wagon Counting

The state railway was having trouble maintaining a regional freight service. A colleague and I were asked to do a review. Examining the historic data, we tried to analyse train tonnages as a function of the recorded train length (reported as the number of wagons) and found the correlation was poor as seen in Figure 2.




Figure 2: Train tonnage and length

On examining the plot of train tonnage against train length we found the data could be much better fitted to two straight lines, one with twice the slope of the other, as in Figure 3.

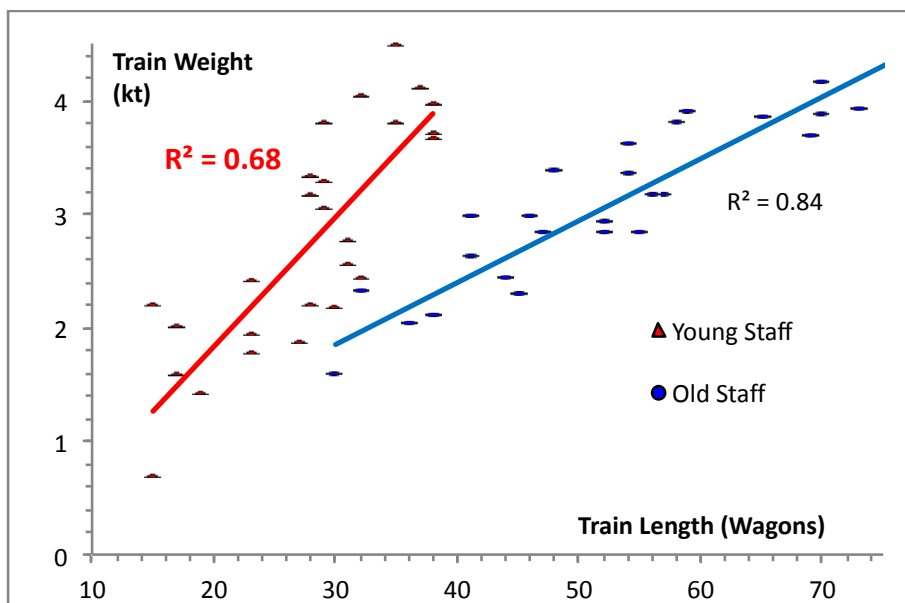


Figure 3: Train tonnage, length and age of staff member

Management could not explain this finding. On talking to a number of staff, the cause eventually became apparent. For their record keeping each day, a staff member went out and counted the number of wagons of each train. About twenty years previously, short wagons had been phased

out and gradually replaced by wagons twice the length, with dual bogeys. If you were over forty you would count {1,2}, {3,4}, {5,6}..., two for each of the new double length wagons. The younger staff had never seen the phased out old-style wagons and would therefore count the new wagons as {1}, {2}, {3}... So the same train could plot on the steeper or on the shallower line, depending on the age and length of service of whoever had recorded the data. Surprisingly, the two cultures of counting existed in parallel, with neither being aware of the other. Correcting the data, we found that train tonnage and length were now strongly correlated, and our analysis could proceed.

Cigarette Production Simulation

My head of department had a call from a tobacco company in Indonesia. Their staff had been writing a simulation model in the computer language “Extend”. They were finding that their model took an inexorable time to run and had heard that we had some experience of programming in Extend. When I got there and had a look at their program, I found that they were modelling the operation of the machine that made filter tips. Each of the hundreds of thousands of filter tips was being programmed as a discrete event, so a day’s production took more than a day of computer time. Changing the paradigm, to consider the filter tips production as a continuous flow instead of a series of discrete events, made the program run in a few minutes, without any loss of accuracy.

Treating the manufacture of each individual filter tip as a discrete event instead of as part of a flow was analogous to using Einstein’s Relativity theory instead of Newton’s Laws to calculate the path of a ball rolling down an inclined plane.

Hospital Laundry Costing

Our state government ran a Hospital Linen and Laundry Service (HLLS) doing the laundry for a number of government institutions, including hospitals, prisons, and retirement homes. In the 1980s there was a push for accountability, and the HLLS was required to charge the institutions an appropriate amount for each item laundered.

Our problem was to estimate a unit cost for each type of laundry item, such as nurses’ uniforms, sheets, operating theatre material, and uniforms.

At first sight, this seemed an intractable problem. The total costs for a full day’s operation were available, but each day’s throughput was a mix of many types of laundry. One possibility was to put through only one type of laundry for the whole day. However, this would be very disruptive operationally. It would also be likely to over-estimate costs, because the equipment would not be used efficiently if only one type of laundry were being processed for the whole day.

Fortunately, the fact that a different mix of laundry went through each day made the problem solvable. Assuming a linear model, the total cost for a day was the summed product of the number of each type of laundry multiplied by the unit cost for that laundry type.

A multilinear regression model was constructed, of the form $Y_n = \sum k_j X_{jn} + e_n$. In our model, Y_n is the total cost of operation in day n , X_{jn} is the quantity of laundry of type j processed on day n . The unit costs k_j for each type of laundry were found, minimizing $\sum e_n^2$. This method uses the standard ordinary least squares (OLS) regression method, except that the intercept is zero (since the fixed costs needed to be apportioned to the laundry types) and the dependent Y “variable” is actually a constant. Having a constant Y_n value does not matter, provided the X_{jn} quantities vary from day to day.

Using this least squares regression we were able to estimate the k_j coefficients and therefore the unit cost for each type of laundry (Everett & McLeod, 1991).

Ore Grade Prediction

Some years later, I found myself consulting for an iron ore mining company. Ore is crushed to yield two products: “lump” and “fines”. Ore had to be selected each day from a range of available sources. The selection had to yield lump and fines products each of specified grade. Each of the two products (lump and fines) had a target grade, not only for iron but also for each of the contaminants: silica, alumina, and phosphorus (Everett, Howard, & Jupp, 2009).

For the ore from each available source, its ore type was known, and an estimate of its total (“head”) grade was also available. However, the head grade samples tended to give biased assays. Also, different types or sources of ore, when crushed, split into different lump and fines grades, with the split varying between the types or sources of ore. Previously, “Run Of Mine” (ROM) trials had been carried out, putting only one type of ore through for the whole shift. ROM trials were costly, of little statistical power (since only a few ROM trials could be afforded) and perhaps distorted the plant performance. Being so costly, the ROM trials were not updated, so did not reflect long-term changes.

On reflection, I realized that the problem was essentially the same as the hospital laundry costing problem, and that the ROM trials were equivalent to devoting the laundry operation to a single laundry type for a whole day.

The problem was solved by constructing a regression model similar to the HLLS model, though somewhat more complex. The predicted Y variable was now multidimensional, being the grade vectors for the daily production of the two, lump and fines, products. The predictor X variables were constructed from the input ore head grades and types.

The regression model enabled lump and fines product grades (output from the crusher) to be predicted from the input ore types and head grades. This helped in the daily choice of ore to generate product close to target grade.

Chances May Not Multiply

At the time of the First Gulf War in 1990-91, the Kuwait Foreign Petroleum Exploration Company (KUFPEC) had operations in a number of areas around the world, including Western Australia. As a fairly centralised organisation, each operation had to send its data back to Kuwait for processing. Their software in Kuwait was routinely backed up to multiple locations, with the bank, the CEO, the main office, etcetera. Since all these backup locations were in Kuwait, they were all lost at the time of the invasion. I had a very nice job reverse engineering their previous year’s reports and data so that the local Perth office could generate its annual reports.

This example illustrates the tendency to assume that, if two safeguards each have a one in a thousand chance of failing, there is only a one in a million chance of both failing. The conclusion is true only if the safeguards are independent. In the case of the Iraq invasion of Kuwait, the safeguards were clearly not independent. Even though you have a Plan B and a Plan C, the same disaster that destroys Plan A can also destroy Plan B and Plan C.

Ore Quality and Opportunity Cost

Iron ore is sold to a contract specification, where the grade of the delivered ore must at least reach a specified percentage in iron and must not exceed specified percentages in the contaminants silica, alumina, and phosphorus. In my earlier days of consulting in this industry, I frequently met staff who were very concerned that the product should meet these specifications but were not concerned if the product exceeded specification. They were not familiar with the concept of “opportunity cost”, the fact that excessively rich ore could have been blended with poor quality ore

that could not otherwise have been sold. As a colleague succinctly put it, “The objective is to please the customers, but not to delight them.”

Ship Loading, Responding to Error

If you aim a gun at a target, and then adjust the aim by the amount you were off centre, and do this repeatedly, your divergence from target will get larger and larger if the error is random. We found that this phenomenon was occurring during ship loading of iron ore for export (Everett, Howard & Kamperman, 2002).

Each ship loading has a plan as to which source stockpiles ore should be taken from. During ship loading, samples were assayed and the resulting assays were used to drive changes in the planned choice of source stockpiles. As a result of this apparent correction, the loaded ships were reported as each being very close to target grade. However, when the ships were unloaded and assayed by the customers, they reported a considerably higher variability than was reported by the producer during loading.

After some discussion, fine tuning was abandoned, and ships were now always loaded according to their original plan. The assays from samples taken during loading were thenceforth used solely for reporting purposes, no longer for control. The reported variability on loading slightly increased, but it now agreed in magnitude with the customers’ reported variability, which had now decreased. The reason was that any assay is subject to random error, and the earlier control system had been responding to this random error, just like a shooter adjusting aim in response to random deviation from target.

The hardest part of the reform was persuading the cargo controllers, who had previously had to be woken up during the night to authorise changes to the loading plan, and who now had uninterrupted sleep. Again, the psychological adjustment was the most difficult part of the reform.

Ore Selection

When a mine is being planned, assay data from exploration drill holes are interpolated to construct a rectangular block model, with each block centroid being located on a regular rectangular grid. Each block has an estimated grade, for the each of the analytes of interest.

It is necessary to distinguish which blocks should be identified as ore and which as waste. If there were only one analyte of interest, such as gold, then the distinction could be made at some cut-off value for that analyte, where the marginal cost of extraction just matched the marginal value of the extracted analyte. Blocks above this cut-off would be accepted as ore, and blocks below the cut-off rejected as waste.

In mining iron ore, the criterion is more complex, since the analyte of value, iron, is accompanied by contaminant analytes, such as silica, alumina, and phosphorus.

The procedures used tend to be commercially confidential, but it appears that the generally used criterion is to apply a cut-off grade to each analyte, such that a block is selected as ore if its assays lie in a multidimensional quadrant, being above a cut-off grade in iron and below cut-off grades in each of the contaminants.

Figure 4 shows that the quadrant cut-off criterion is inconsistent: all the blocks mapping below the slanting line but outside the quadrant would be excluded by the criterion. If they were combined to a single block they would map as the large blob well within the quadrant, so would be accepted.

It can be shown that a composite cut-off criterion (as shown in Figure 4) will yield the maximum tonnage of ore at a specified grade vector. The composite cut-off function is the sum of each ana-

lyte grade multiplied by a coefficient, where the coefficients are positive for iron and negative for each of the contaminants. A block is selected as ore if its composite function score is above some cut-off value.




Figure 4: Ore selection by quadrant and composite cut-off criteria

Iterative adjustment of the coefficients and of the cut-off value identifies the maximum tonnage that can be extracted for a nominated product grade vector. The coefficients can be interpreted as the marginal value or cost of each analyte.

The composite cut-off criterion will always yield at least as much ore as the quadrant criterion, and generally five to ten percent more. Since deposits are typically a few hundred million tonnes, and iron ore has a value now well in excess of \$100 per tonne, it can be seen that an extra five or ten percent of recoverable ore is not trivial.

I presented this analysis to industry practitioners at a recent geometallurgical conference (Everett, 2011). The session response and conversations afterwards confirmed that the quadrant cut-off criterion is the predominant method for ore selection.

Prior to the general availability of computer power, it would have been very difficult to construct and apply an appropriate composite cut-off criterion, and this probably explains why the quadrant criterion was originally adopted.

The continuance of the quadrant cut-off criterion provides a convincing example of paradigm paralysis, where an inappropriate paradigm has survived although the technology has been available for some time to enable the more efficient composite cut-off criterion.

Two more factors may have contributed to paradigm paralysis in this example. Proprietary data processing packages specific to the industry are used as “black boxes”, without their workings and assumptions being understood or even accessible to the users.

Secondly, the mining industry is very subject to secrecy, so data analysis methods tend not to be subjected to the scrutiny that would come with more open discussion. This secrecy is understandable, since there is considerable commercial advantage in having better methods than one’s competitors, but can be counterproductive if the lack of discussion means everyone’s method is sub optimal. Because staff tend to move frequently between operating companies, the perceived benefits of secrecy are negated while the disadvantages remain. Professional conferences provide one opportunity for information and methods to be discussed and shared. Unfortunately, at least in the mining industry in Australia, the larger companies have over recent years become more reluctant for their procedures to be disclosed at conferences.

Batch versus Process

One paradigm shift that I have observed in the iron ore mining industry is the transition from batch to process thinking (Everett, 2001). In earlier times (and still in some operations) the emphasis was on building individual stockpiles or on loading individual ships to target grade. This tends to lead to “hunting” behaviour, where solving today’s problems generates tomorrow’s problems, because the corrective action is too late or too large. It also encourages sub-optimisation, where each section of the organisation concentrates on making its own performance look good. For example, the port operations blame the mine when things go wrong, and vice versa.

At a further cost, batch thinking leads to excessive stockpiling, with multiple stockpiles low or high in each of the analytes being kept for use in emergencies. This batch production, controlling the grade of a series of discrete stockpiles, has now been replaced by controlling the exponentially smoothed grade, where the grade is exponentially smoothed over a “half-life” corresponding to the typical stockpile tonnage, eliminating the discontinuities that previously occurred at the start and end of each stockpile build.

Constraints versus Objective Functions

A separate but related issue that affects whether a system runs smoothly and continuously is the question as to whether control should be by constraints or by contribution to an objective function. For example, iron ore grade needs to be within set limits not only for its iron content but also for silica, alumina, and phosphorus. The control can be applied by constraints, requiring that the grade must not lie outside definite upper and lower bounds. This leads to discontinuities: no one worries as the grade approaches a limit, then suddenly all hell breaks loose when the constraint is breached. The problem is increased by the fact that any assays on which decisions are based are themselves subject to measurement error, so the system may appear to be outside the control limits when actually it is not, and vice versa.

Because of the human tendency to confirmation bias, the use of constraints may leave attention unduly focussed on one problem (such as alumina content) when another (such as silica) is creeping up to bite us. A smoother operation can be achieved by replacing the discontinuous constraints by contributors to a continuous objective function, as illustrated in Figure 5.

In the iron ore example, the constraints for each analyte were replaced by an overall objective function. For each analyte (iron, silica, alumina, or phosphorus) a component stress was defined as the grade minus the target, divided by the tolerance. This stress measurement (positive or negative) replaces the previous constraints. The dimensionless component stresses are each squared, and summed to give a “total stress”. The total stress is used as a single overall objective function, to be minimised.




Figure 5: Constraints versus an objective function

Figure 5 shows how the quadratic objective function provides much smoother control than the discontinuous constraints. Being comprised of contributions from all the analytes, suitably weighted by their tolerances, minimising the objective function automatically controls all the analytes simultaneously.

Replacing the discontinuous constraints by a smooth objective function recognises that the process being controlled is a statistical process, which will always have a small probability of generating occasional outcomes beyond the control limits. In the iron ore example, this can be accommodated by occasionally paying contractual penalties, whose cost amounts to less than the cost of having an over-controlled system.

Error Estimation

Assays of iron ore samples are subject to error, comprising three separate contributions, due to errors in the sampling, preparation, and assaying procedures. The International Standards Association publishes a manual for estimating these errors from a series of duplicate measurements (ISO, 2002), and this is used as the industry base for establishing assay reliability.

The published ISO procedure uses “mean absolute deviations” instead of “root mean square deviations” to establish variances, thus making the implicit assumption that measurement errors are normally distributed, which they demonstrably are not. This distinction is important because, provided their distributions are independent or uncorrelated, variances add but mean absolute deviations cannot be compounded (even when squared) unless the underlying distributions are also normal.

Use of mean absolute deviations may have been justifiable before the availability of computer power to calculate root mean square deviations, but that is clearly no longer an excuse for paradigm paralysis. The ISO procedure also prescribes the discarding of outlier results, without searching for and removing the causes of the outliers. The net effect is systematically to generate overly optimistic estimates of precision. (Everett, Howard, & Beven, 2011).

A Tentative Taxonomy of Paradigm Problems

For the examples discussed above, it is worth considering a tentative taxonomy of the paradigm problems encountered and the approaches that were used to overcome resistance to reforming them. For this classification, I am indebted to the editorial input of Grandon Gill.

The types of paradigm problem identified below are not mutually exclusive, but can actually reinforce each other, with any one situation involving a mixture of paradigm problems.

Alternative Paradigms in Small Worlds

In any organisation, old-timers and newcomers tend to work and socialize in separate groups, with little communication between the groups.

Examples of this occurred in many of the examples. In “cardboard box cutting”, the operator minimising the width rather than proportion of wastage had clearly not discussed the method with any of his colleagues. Similarly, in “railway wagon counting” there were two different age groups of employees whose different paradigms for wagon counting were each self-consistent but mutually inconsistent.

In several of the mining examples, this problem arises because of industrial secrecy, where operators doing similar work in different companies are not able to discuss their work with each other. Here, consultants and conferences can both help the flow of information and paradigm revision

Solutions to this problem may lie in a freer interchange of information within the organization. For example, when the two groups of wagon counters learned that they were using different methods, management was able to prescribe a future policy and the problem disappeared. In that case, there was no psychological problem, only a misunderstanding to clear up.

With the cardboard box cutting, there was a psychological problem, since the previous decision maker stood to lose face. Fortunately, as described, the change was made tactfully without hurting the man’s feelings.

Problem Representation Inadequate

The underlying model being used may be totally correct, but inappropriate or incomplete for efficient decision making.

In the “cigarette production simulation”, the production of each individual filter tip is indeed a discrete event. However, a simulation model treating each filter tip as a discrete event takes an enormous time to run. Modelling the filter tip production as a flow is slightly less realistic, but makes the simulation feasible without loss of accuracy.

In the “hospital laundry costing” and “ore grade prediction” it had been thought that production would have to be run for a day or a shift using only one type of input to estimate the cost or performance parameters for that input. Formalizing the problem as a linear regression model enabled parameters for each type to be extracted from a varying daily mix.

Invalid or Inefficient Heuristic

Sometimes the decision makers may be using a paradigm which, while superficially reasonable, is actually wrong. In the KUFPEC example, it had been implicitly assumed that probabilities multiply, so that multiple backups would give adequate security. However, probabilities do not multiply unless events are independent: in the case of Kuwait, the one invasion lost all the backups simultaneously.

Being a one-off event, there was no solution, except by rebuilding the software (although technology has now to some extent removed the problem, with data in the “cloud” being stored in multiple diverse locations). The problem for the company became an opportunity for the consultant who had the job of reverse engineering the company’s software.

“Ship loading” provided another example of an invalid but superficially reasonable model: adjusting the grade to bring the measured grade back to target is invalid if the measurements include error. The solution here was by education through analogy: working through the gun aiming example convinced the operators, who changed their practice. This situation again contained a psychological element, since those hardest to convince were those whose work was most inconvenienced, but also made significant, by the former practice of continual adjustment.

The erroneous “ore selection” and “error estimation” paradigms both probably had their origin in the days when the computer ability for a more valid model did not yet exist. In the days of manual computation it was easier to apply cut-off grades than to work with a composite function. Similarly it was easier to compute “mean absolute deviations” than “root mean square deviations” in the days before computer power made the computations trivial. In both cases, inertia (and lack of communication) delayed the adoption of more valid models. Paradigm shift was slowly achieved, through demonstration of the benefit, using clients’ real data.

Inappropriate Objective Function and/or Constraints

Many procedures require the optimisation of an objective function subject to constraints. If the production is ongoing, this amounts to a control system. A control system works more smoothly if it does not involve discontinuities. Replacing a series of batch productions by the control of the exponentially smoothed grade for an equivalent ongoing process can remove many such discontinuities (even when the physical production continues to be a series of batches).

Similarly, replacing constraints by adding equivalent terms to the objective function can replace discontinuities by a continuously changing pressure, and thus aid control of the system as a smooth statistical process.

Conclusion

The examples discussed illustrate that paradigm paralysis is frequently encountered in ongoing organisations, particularly in the treatment of data and extraction of information to enable informed decisions. The information sector is particularly vulnerable because, prior to the advent of cheap and accessible computing power, many sub-optimal procedures were followed that are no longer necessary.

The examples given show that there is plenty of opportunity for academic consultants to provide useful inputs, even to areas outside their own specific area of expertise. It should be emphasised that for such work to be successful the outsider must respect and learn from the domain knowledge of the clients, work with them rather than for them, and be cognisant of the psychological barriers that occur in implementing paradigm shift.

References

- Bullard, E. C., Everett, J. E., & Smith, A. G. (1965). The fit of the continents around the Atlantic. *Philosophical Transactions of the Royal Society*, 258, 41-51.
- Everett, J. E. (2001). Iron ore production scheduling to improve product quality. *European Journal of Operational Research*, 129, 355-361.
- Everett, J. E. (2011). Ore selection and sequencing. *Proceedings of Geomet 2011, AusIMM*, Brisbane, Sept 2011, 197-204.

Paradigm Lost

- Everett, J. E., Gordon, J., & Jarvis, G. (1976). Trimming the trim: An heuristic approach to minimizing paper wastage in the manufacture of cardboard boxes. *Economic Activity*, 19, 42-45.
- Everett, J. E., Howard, T. J., & Beven, B. J. (2011). Precision analysis of iron ore sampling, preparation and measurement overcoming deficiencies in the current standard ISO 3085 product. *Mining Technology*, 120(2), 65-73.
- Everett, J. E., Howard, T. J., & Jupp, K. (2009). Predicting lump and fines finished product grades and lump percentage from head grade. *Mining Technology*, 118(2), 102-108.
- Everett, J. E., Howard, T. J., & Kamperman, M. (2002). Information for decision support, information for performance evaluation – But don't mix your drinks. Proceedings of the *Informing Science and Information Technology (I²SITE) conference*, June 2002, Cork, 427-439.
- Everett, J. E., & McLeod, P. (1991). Direct labour cost allocation for a multi-product hospital laundry. *Accounting and Finance*, 31(1), 39-51.
- Gill, T. G. (2010). *Informing business: Research and education on a rugged landscape*. Santa Rosa, California: Informing Science Press.
- ISO 3085. (2002). *Iron ores - Experimental methods for checking the precision of sampling, sample preparation and measurement* (4th ed.). International Standards, (SAI Global: Sydney).
- Kuhn, T. S. (1996). *The structure of scientific revolutions* (3rd ed.). Chicago and London: Univ. of Chicago Press.
- Plous, S. (1993). *The psychology of judgment and decision making* (p. 233). McGraw-Hill.

Biography



Jim Everett began as a research geophysicist at Cambridge and the Australian National University. His family always said “what are you going to do when you grow up”, so he then joined industry, as a petroleum exploration geophysicist. With a computer background, he was put onto project evaluations. Losing arguments with accountants, he did a couple of Economics and Commerce degrees at the University of Western Australia (UWA), and found the accountants were usually right. UWA was starting its MBA course and wanted people with industry experience to teach on it. With a young family, fieldwork was now less attractive, so Jim jumped at the chance of returning to academia, spending thirty years in the UWA Business School.

Since retirement as Emeritus Professor of Information Management, he divides his time between research and consulting to Western Australia's booming mining industry. Most old folk lose their faculties, but he has recently gained one, joining the UWA Earth and Environment Faculty as an Adjunct Professor in the Centre for Exploration Targeting.