ORiON, Vol. 3, No. 2, pp. 85-100 *ISSN 0259-191X*

THE MYSTERIOUS CASE OF THE MISSING HALF-TON OF GOLD

Douglas M Hawkins
Department of Applied Statistics, University of Minnesota
St Paul, MN 55108, U S A

and

Hugh E Bartlett
Johannesburg Consolidated Investments
Box 590, JOHANNESBURG, 2000 South Africa,

Summary

A milling circuit in the concentrator processing ore at one of the gold mines managed by Johannesburg Consolidated Investments had apparently been losing a substantial part of the gold input to it for several months. At the time of investigation, the gold unaccounted for totalled nearly half a ton, worth some $5 million. A specially adapted distributed lag model was fitted to measurements made of the input and output, and far from confirming the apparent loss of gold, indicated that the plant's output balanced its input. The apparent contradiction between the total input and output and the more detailed accounting of the distributed lag model (which included variable time lags explicitly) gave rise to a closer investigation of the circuit which led to an explanation of the apparent discrepancies, and the actual location of a large amount of gold.

Prologue

During 1984, the staff of a gold mine with which the authors are concerned noticed an alarming and initially inexplicable trend in the balance data of one of the plants processing ore. For four successive months, the monthly total amount of gold in the ore fed into one of its plants greatly exceeded the amount of gold metal obtained at the end of the process, and at the stage at which the authors became involved a total of nearly 500Kg of gold, worth some US$5 million, was unaccounted for.

The figures of the monthly input and output measured in grams of pure gold were as follows:-

Month	Daily mean input	Daily mean output	Shortfall
1	15891	14287	1604
2	19435	14809	4624
3	27277	22447	4831
4	30937	26116	4816

The shortfall amounted to 17% of the input. There was no very obvious explanation for this large amount of missing gold, and management was understandably concerned. The ever present possibility of a theft was being taken seriously, and a wide-ranging investigation into the possible leak was under way. As part of this exercise, a detailed statistical analysis of routine plant records was undertaken, and ultimately gave rise to the results reported here.

The plant and process

As most readers are probably unfamiliar with the metallurgy of gold extraction, a very brief thumbnail sketch of the operation of a gold concentrator may be in order. The gold-bearing ore is blasted out underground in the form of large rocks of diameter up to 200 mm which are then brought to surface. The ore is sampled by stopped-belt sampling, in which at regular intervals the conveyor belt carrying the ore is stopped; a "former" (essentially a pair of parallel boards half a metre apart) is placed over the belt to demarcate a sampling interval; and all material within the former is swept off the belt into a sample bucket for chemical assay. The assay values so obtained, together with the throughput tonnages, constituted the "input" portion of the data used in this exercise. The ore passes into a rotary mill (on some mines having first been put through a jaw crusher to reduce the particle size to some extent) for reduction to very fine particle size.

Mills come in many different shapes and sizes, but those used in gold mines are typically some 3 m in diameter and 5 m long and rotate at a speed of the order of one revolution per second. Enough ore and water (and in some designs of mill, steel balls) to keep the mill about half full is fed continuously into one end of the mill. Friction against the sides of the mill keeps the charge in motion, so that as the mill rotates, it is dropped from near the top position onto the bottom and the impact of the charge against itself grinds the particles of ore down. This milled material passes down the mill and out the far end, typically through or near the axis. Here a cyclone separates the finer particles from the coarser. The latter are recycled into the feed end of the mill with the water supply for further milling. (The actual milling circuit of the plant under discussion was of this broad type, but involved several serial stages of milling, as well as preliminary crushing of the ore prior to any milling). The fine ore pulp finally emerging from the milling is first "thickened", by being put into large tanks, where the solids settle out and the surplus water is drawn off. The solids are then transferred to air-agitated tanks to which cyanide is added to dissolve the gold.

After this dissolution process, the solids are filtered off, and may be further processed to extract useful products such as uranium and sulphur, while the gold-bearing solution is de-aerated and zinc dust is added. This reprecipitates the gold in metallic form, and it is filtered off for several further refining steps before being smelted into gold bars.

The feed to the cyanidation in the plant under discussion is sampled and assayed regularly, and this provided the figures earlier given for the plant output. Thus the material balance being attempted was from the stage of the start of the ore processing to the feed for cyanidation.

Rough volumetric calculations showed that the total capacity of the processes between these two points amounted to about four days' throughput, and this suggested that when one was balancing over a period as long as a month, the fluctuations caused by variations in the amount of ore in process should be negligible.

While the plant had been running continuously for some years and could therefore be supposed to have settled into a steady state, at about the start of the period of severe imbalance its feed had been switched from a relatively poor material (of the order of 1 gram of gold per ton of ore) to an ore about five times richer. The changeover to the richer ore did not involve any special shutting down of the plant, and so was not expected to have any appreciable effect on the steady state - a perception that, as we shall see below, was quite unfounded.

The statistical analysis undertaken

Data on the input and output were available on a daily basis for a period of four months. After the usual data verification and computation of descriptive statistics, the first analysis made was of the cross correlation of the output with the input at various lags back in time. The auto- and cross-correlations at lags up to one month are sketched in Table 1. As pointed out by a referee, the process was not time-stationary over the period in which the data were gathered. While this does not invalidate the subsequent regression analysis, the reader should be warned that the nonstationarity has the effect of reducing the correlations at long lags from the values that would be obtained from the same process after it had reached steady state - in particular, at very long lags the nonstationarity induces negative correlations where the steady-state correlations would be zero.

There are several noteworthy correlations. One is the strong weekly input cycle, due to the fact that while the plant itself runs continuously, the crushers are closed on Sundays. Also striking is the large input correlation of 0.61 at lag 28 days, where the weekly cycle is in phase with the organizational one of tidying up for the month end. While the gold loss is considerably more random, this also shows the substantial month-end effect. The output is strongly autocorrelated. This is to be expected, since the operations have the effect of smoothing fluctuations in input, leading to a more even output stream.

Table 1. Correlations at various lags

Correlating:	Output	Output	Output	Input	Input	Loss
with lagged:	Output	Input	Loss	Input	Loss	Loss
Lag (days)						
0	1.00	0.24	0.36	1.00	-0.83	1.00
1	0.66	0.09	0.29	0.26	-0.07	0.23
2	0.49	0.16	0.12	-0.06	0.26	-0.18
3	0.37	0.21	0.01	0.00	0.16	-0.15
4	0.36	0.24	-0.03	0.08	0.08	-0.10
5	0.27	0.39	-0.23	0.00	0.18	-0.31
6	0.32	0.37	-0.18	0.10	0.04	-0.15
7	0.36	0.35	-0.13	0.56	-0.35	0.26
8	0.40	0.12	0.10	0.24	-0.10	0.16
9	0.38	0.11	0.10	0.05	0.11	-0.04
10	0.32	0.15	0.03	0.16	-0.05	0.07
11	0.29	0.26	-0.10	0.19	-0.04	-0.02
12	0.27	0.34	-0.17	-0.02	0.18	-0.29
13	0.25	0.34	-0.18	0.08	0.07	-0.18
14	0.35	0.34	-0.13	0.53	-0.30	0.22
15	0.39	0.25	-0.03	0.25	-0.04	0.03
16	0.34	0.07	0.12	-0.03	0.24	-0.19
17	0.27	0.10	0.05	0.02	0.07	-0.05
18	0.32	0.28	-0.09	0.20	-0.15	0.11
19	0.39	0.24	-0.02	0.11	-0.03	0.02
20	0.34	0.26	-0.06	0.08	0.06	-0.10
21	0.30	0.37	-0.19	0.54	-0.33	0.24
22	0.31	0.25	-0.07	0.17	0.00	-0.05
23	0.34	0.13	0.07	-0.05	0.18	-0.16
24	0.28	0.18	-0.02	0.06	0.04	-0.05
25	0.34	0.34	-0.13	0.20	-0.14	0.08
26	0.41	0.35	-0.11	0.09	-0.02	-0.04
27	0.38	0.13	0.08	0.12	-0.07	0.12
28	0.35	0.11	0.08	0.61	-0.40	0.45
29	0.31	0.15	0.03	0.17	0.03	-0.01
30	0.30	0.13	0.04	-0.10	0.32	-0.29

The next stage of analysis that suggested itself was an attempt to model the residence time of a particle of gold in the system. For this, a distributed lag model was the obvious choice.

While distributed lag models are widespread in economic work, and have a distinguished pedigree going back to R A Fisher (1924), we felt it advisable to use a non-standard model taking specific account of facts known from chemistry and from the characteristics of the plant.

Letting Y_t denote the output, and X_t the input to the plant on day t, equations of balance would lead one to equations of the form

$$Y_t = \sum_{k=0}^{\infty} a_k X_{t-k}$$

where the coefficients a_k reflect the distribution of the residence times. A distributed lag model consists of writing a_k as a function of k and perhaps of other unspecified parameters that have to be estimated. For example, the polynomial distributed lag model holds that

$$a_k = \sum_{\ell=0}^{L} b_\ell \, k^\ell$$

over some range, a_k being zero outside that range. For our application, rather than using the general-purpose polynomial model, whose applicability to the plant is not apparent, it is far more natural to invoke results of queueing theory, or their formal chemical equivalents. For example, if the plant constituted a single well-mixed vessel, then the residence time of a particle would follow an exponential distribution, so that

$$a_k = e^{-\beta k}$$

where the parameter β is simply the ratio of the flow rate per unit time to the total volume of the vessel. As a more complicated model, if the process were ℓ perfectly mixed vessels in series, then (by the familiar Erlang multistage argument) the residence time would be

$$a_k = \frac{(\beta k)^{\ell-1}}{(\ell-1)!\beta^{\ell}} e^{-\beta k}$$

More generally, if the process were regarded as some mixture of up to L perfectly mixed stages, then the residence time would be modelled as

$$a_k = e^{-\beta k} \sum_{\ell=0}^{L} c_{\ell} (\beta k)^{\ell}$$

where the c_{ℓ} are as-yet-unspecified coefficients.

In terms of the original data, this model can be written:-

$$Y_t = \sum_{\ell=0}^{L} \sum_{k=0}^{\infty} X_{t-k} e^{-\beta k}[c_{\ell} (\beta k)^{\ell}]$$

Gathering together the terms in ℓ, this can be seen as a linear regression model predicting Y_t from the constructed predictor variables

$$\sum_{k=0}^{\infty} X_{t-k} e^{-\beta k}(\beta k)^{\ell}, \qquad \ell = 1,\ldots,L$$

Despite the appearance of gamma-like terms, the above model should not be confused with the "gamma distributed lag" discussed, for example, in Schmidt (1974). The latter model involves an unknown power, as well as an unspecified value of the "half-life" parameter β.

One final analytical step was performed. The main by-product anticipated from this analysis was a formal test for the conservation of the gold in the system $(\Sigma_0^\infty a_k=1)$. To facilitate this, the problem was rewritten in terms of orthogonal polynomials $P_\ell(k)$ rather than powers of k as:-

$$Y_t = \sum_{\ell=0}^{L} d_\ell \sum_{k=0}^{\infty} X_{t-k}\, e^{-\beta k} P_\ell(k)$$

where the polynomials $P_\ell(k)$ (whose construction we discuss below) were chosen to be orthogonal with the weighting function $e^{-2\beta k}$. In principle, the coefficients of these polynomials are "well known", a closed formula being given in Gottlieb (1938). In practice, however, since the actual domain used was truncated at a finite previous time (K=30), we found it easier to obtain the necessary values numerically using a direct weighted Gram Schmidt orthonormalization of the sequence of powers of k.

This orthonormalisation proceeded as follows. First, we note that we do not need the explicit representation of the polynomials - all that is needed is a table for the range of k and ℓ values used of the weights $e^{-\beta k}P_\ell(k)$ taking the sequence X_t into the constructed variables. To get these, we first set up the array of weighted powers $k^\ell e^{-\beta k}$ $k=0,\ldots,30$, $\ell=1,\ldots,L$. The first row of this is rescaled to sum to 1, thereby defining the weight table $e^{-\beta k}P_1(k)$ for $k=0,\ldots,30$.

Having found the first weight column, we then regress each subsequent column on it, and subtract out the regression, leaving a residual which is by definition orthogonal to the first weight column.

Rescaling the second column to unit length gives the second column of the weight table, and we orthogonalise the 3^{rd} ... columns to it by regressing each of them in turn on it, and replacing each column by its regression residual.

The orthogonalisation thus consists of a sequence of simple linear regressions and is perhaps easier to do than to describe. As we find each column of the weight table, we remove its regression contribution from each column to the right of it. To complete the orthogonalisation we just normalise each column to unit length.

Th orthogonalization had the practically useful and important implication that the first constructed variable is a weighted average of the inputs up to that time, while the subsequent terms are all "contrasts", averaging to zero. This meant that for our application the sum of the coefficients $\Sigma_k a_k = d_0$, and so the test for the conservation of gold reduced to the simple issue of testing whether d_0 was 1.

Fitting and results

The autocorrelation between the different Y_t, while appreciable, was not so high that we felt it necessary to use a generalized least squares approach. Instead, the model was fitted by ordinary least squares using the REGPAC regression package (Galpin 1981). The first month's output data were ignored, but the corresponding input data were used for the lagged terms of the model. The results of Baltagi and Ferry (1985) suggest that the truncation of the terms at lags greater than 30 days, as occurred for the early part of the terms at lags greater than 30 days, as occurred for the early part of the output data, should not have a severe effect on the quality of the resulting fit.

Successive orthogonalised polynomials were fitted to establish the order of process best describing the data. The nonlinear parameter was taken to be the 0.25 days suggested volumetrically, though reruns with different values gave much the same final conclusions.

Only the coefficients for lags 0 and 1 were significantly different from zero. Their estimated values and standard errors were:-

Term	Estimated value	Standard error
d_0	1.027	0.014
d_1	0.174	0.076

As a consequence of the orthonormalisation of the polynomials, this means that the point estimate of the fraction of the gold fed into the plant which ultimately emerges is 102.7%, with a standard error of 1.4%, indicating that the very severe underaccounting found previously has been replaced by a slight (and statistically not significant) overaccounting.

It is striking that this point estimate of 102.7% accounting is very (and statistically significantly) different from the point estimate of 83% obtained directly as the ratio of each month's shortfall to its input. The difference between these two point estimates of percentage conservation is a measure of the distorting effect that the time lags between input and output have on the plant balances obtained from "snapshot" cross-sectional data such as the

The coefficient d_1, again thanks to the orthonormalization, conveys no information about the conservation of mass within the plant, but it does imply that the plant does not conform to the simple "first order reaction" used as an initial working model, and that it in fact has considerably longer lags before the material emerges. It is of interest to transform the coefficients d_0 and d_1 back to the original scale of the lag distribution coefficients a_k - these are

Lag k	0	1	1	3	4	5	6	7	8	9	10	>10
$100a_k$	17	17	15	13	12	10	8	7	6	4	3	7

This shows that, while most of the material has quite a short residence time in the system, 7% is in process for more than 10 days.

It has been pointed out by the referee that some algebra produces a direct estimate of the amount of gold in the system at time t. This is:-

$$\sum_{k=1}^{K} a_k \left[\sum_{s=0}^{k-1} X_{t-s} \right]$$

which relates the not-yet-accounted gold at any instant to moving totals of the historical sequence of inputs, weighted by the residence time coefficients a_k.

As regards the primary motivation for carrying out the analysis - that of determining whether there was any statistical evidence of a systematic loss of gold - the table is really quite clear in saying that d_0 is close to 1 in relation to its standard error, and therefore that the data give no reason to reject the hypothesis of conservation of the gold entering the system. While this is encouraging, it should be noted that the 95% confidence interval is fairly wide, covering $\pm$ 3% of the gold input, an amount whose dollar value might be a source of some discomfort to management. This concern may of course be resolved by simply extending the data gathering and analysis exercise into either the previous or the subsequent months' data.

Epilogue

The statistical analysis had produced the superficially highly implausible conclusion that, despite the apparent four-day residence time and the persistent and large imbalance of input and output for four months, there was no evidence of the system's being bled of its gold. Even at its high density, half a ton of gold has an appreciable volume (25 liters in pure metal form), and there really was only one place such an amount could possibly be: in the mills. Thus the plant was stopped some time later, and the sealed cover was taken off one of the mills. On inspection, there was visible in the mill a pile of entrapped fine material, which, when sampled and assayed, was found to be as much as 2000 times richer than the ore feeding the plant. Samples of material from different parts of the mill were assayed, and these assays were weighted up by the amount of material in each part to get a ball-park estimate of the amount of gold in the mill. This rough estimate came to a quarter of a ton of gold. While not providing the entire half-ton that was unexplained, this clearly confirmed that mills could hold very large amounts of gold, and that, between the mill that was examined and the others that were not, it was quite plausible that all the missing gold could be accounted for.

In retrospect, it was not difficult to trace the problem. Gold is extremely dense (19.5 gram per ml), and the drawoff of fine product from the outlet end of the mill depended on the rotating charge in the mill being carried along the mill, and out through the mill's axis. Due to their density, many particles of gold were falling down too straight to be caught by the lifters' at the mill end, and instead were simply going round and round in the mill. This did not lead to any increase in the total volume of material in the mill, but led to a considerable enrichment of its gold content. Thus, while the system appeared to remain in steady state following the increase in the feed grade, it was in a transient state for a very long period, and the final steady state involved a large additional amount of entrapped gold.

The design used for these mills, in particular as regards the method of output of milled material, is a very common one in the industry. It thus follows that many other gold plants have the same situation of a large and very valuable slug of fine gold being locked up in the bottom of their mills and making no contribution to cash flow. While this fact is by no means unknown, its potential magnitude does not appear to be so well appreciated, and nor does the benefit that could be reaped by using a mill design in which, for example, milled product was drawn from the circumference rather than the centre of the mill. Such a design has indeed been suggested and implemented in the past (Mokken et al 1975), but is not common, largely because of perceived engineering difficulties.

While this discussion has concentrated on the formulation and analysis of the model, it should also be noted that a prerequisite for the analysis pinpointing the discrepancy to a relatively small part of the overall processing operation was the availability of data bracketing the mills. In view of a wide-spread belief that it is not feasible to sample unmilled ore, on many (and perhaps most) mines, no attempt is made to take such samples, and the only balance possible is a very coarse one between gold on the face and the feed to cyanidation. If the only data available had covered so large an the feed to cyanidation. If the only data available had covered so large an expanse of operation, it is doubtful whether the problem could have been traced so quickly to the mills, and this may be a major factor in the apparent lack of a widespread concern in the gold mining industry about the amount of gold locked up unproductively in mills. Thus the success of this project is an endorsement not only of the statistical approaches used, but equally of the management decision to implement material accounting over unit processes and not just over the whole surface plant.

Acknowledgements

The authors are grateful to the management of the Gold Division of Johannesburg Consolidated Investments for permission to publish the results discussed; to Mr A H Pienaar for providing the data; to Dr D S Lubinsky for helpful discussions of some mathematical details of the orthonormalisation, and to Dr D Henderson for helpful comments on some metallurgical considerations relating to the data. The authors further acknowledge gratefully the comments of a referee which cast interesting new light on some aspects of the analysis.

Much of the writing up and development of the ideas of this paper was carried while the first author was Senior Consultant to the National Research Institute for the Mathematical Sciences, CSIR, Pretoria, and the support of NRIMS is also gratefully acknowledged.

References

Baltagi, H., and Ferry, J. R., (1985), "Small sample results on Schmidt's truncation remainder to the gamma distributed lag", Statistische Hefte, 26, 87-95.

Fisher, R. A. (1924), "The influence of rainfall on the yield of wheat at Rothamsted", Philosophical Transactions of the Royal Society of London, B, 216, 89-142.

Galpin, J. S., (1981), Regression Package REGPAC (Version 3), Report SWISK 25, CSIR, Pretoria, South Africa.

Gottlieb, M. J., (1938), "Concerning some polynomials orthogonal on a finite or enumerable set of points", American Journal of Mathematics, 60, 453-458.

Mokken, A. H., Blendulf, G. K. I., and Young, G. J. C., (1975), "A study of the arrangements for pulp discharge on pebble mills, and their influence on mill performance", Journal of the South African Institute of Mining and Metallurgy, 75, 257-289.

Schmidt, P., (1974), "An argument for the usefulness of the gamma distributed lag model", International Economic Review, 14, 120-131.