Mining Information

Author(s): Ralph Penner and Donald G. Watts

Source: *The American Statistician*, Vol. 45, No. 1 (Feb., 1991), pp. 4-9
Published by: Taylor & Francis, Ltd. on behalf of the American Statistical Association
Stable URL: http://www.jstor.org/stable/2685231
Accessed: 06-04-2015 12:55 UTC

RALPH PENNER and DONALD G. WATTS*

A data set concerning two processes for drilling holes in rock is presented in the context of an actual consulting session. The problem is easily understood and has important intrinsic value, and the data set is extremely rich pedagogically.

1. AUGUST 22, 1989

Professor Watts?

Yes.

Hi. My name is Ralph Penner. I'm a third-year mining student working for Professor de Souza.

And what can I do for you, Ralph?

I have a statistics problem which I am sure is simple for you, but I need some advice.

Fine. Tell me about it.

I'm trying to fit curves to two sets of data, and I know that some of the data points are not as reliable as others. I have some plots here. You see these large values indicated by squares . . . (See Fig. 1. The data are listed in Table 1.)

Wait, Ralph. Before we go too far, what's the real problem?

Well, as I said, I'm trying to fit curves to these data . . .

No, no. What's the real *problem?*

Oh. Well, Dr. de Souza is trying to find out whether drilling holes is faster using wet or dry drilling.

Hmmm. And how did the data get generated?

Well, in about a 10-foot by 20-foot area, we drilled six holes, three dry and three wet . . .

What exactly do you mean by "dry" and "wet"?

In a dry hole, we force compressed air down the drill rods in order to flush the cuttings and drive the hammer, and in a wet hole we force water. In each hole we recorded the time, in minutes, to advance 5 feet, to a total depth of about 400 feet.

Why 5 feet?

Well, as the hole gets deeper, you have to add rods to the drill, and each rod is 5 feet long.

So the depths are extremely accurately determined. And what you did then was to measure the time it took between additions of a rod?

Basically, yes. Now if you look at the plots you'll see that the drilling time increases with depth, which could be caused by the greater mass of the drill rod string.

How?

The drill is not only rotated, but is hammered down.

Kind of like a jackhammer. And what does the drill look like?

The bit is like a thick plate 6 inches in diameter, with tungsten carbide "buttons" on the bottom, and with gaps in the plate so the cuttings can ride up above the bit and be expelled from the hole.

I see. And how does the fluid get down there?

The rods are hollow and the bit has two holes on its face, so we just force the fluid down the center, which then flushes the cuttings past the bit and out of the hole.

Makes sense. And is it worthwhile to reduce the drilling time?

Oh yes. If we could increase the penetration rate by 20%, that would provide considerable savings. We've also been looking at bit life, and it appears that the life of bits is increased by drilling dry too.

That's good news. I'd like to see that data sometime.

Sure.

But getting back to the data here, you drilled six holes, three wet and three dry, in a 10- by 20-foot area. How did you lay out the drilling pattern?

We just drilled three dry holes in a line, about 5 feet apart, then moved over about 5 feet and drilled the wet holes beside the dry ones.

In a rectangle. Like the six spots on a die?

Yeah.

I see. It might have been a good idea to "balance" the holes so that you had a wet, a dry, and a wet in one line, and a dry, a wet, and a dry in the other. That way, if you were so unlucky as to have a vertical fault line which ran between to the lines you set up, with soft rock on one side and hard on the other, you would avoid making the mistake of attributing a shorter drilling time to the process which happened to hit the soft rock.

Hmmm. I guess that's right. But in this area, and in such a small region, that's not likely.

I agree, but you might want to keep that kind of consideration in mind the next time you plan an experiment like this. Now you said that parts of the data were more reliable than other parts . . .

Yes. Dry hole 3 was the best. There was no binding, and everything went very smoothly. You can see that from the individual hole plots. But with the other holes, the drill would bind, or a hydraulic hose would rupture, or something else would go wrong, and we would get longer times. For example, with dry hole 1 at 200 feet I recorded in my log book that the drill got stuck and so I know that time isn't right. And notice for dry hole number 2, the drill rods got so badly stuck that we couldn't drill deeper than 300 feet. And we could barely remove the drill rods from the hole!

*Ralph Penner is a student, Department of Mining Engineering, and Donald G. Watts is Professor, Department of Mathematics and Statistics, both at Queen's University, Kingston, Ontario, K7L 3N6 Canada. The authors thank the Mining Industry Research Organization of Canada for permission to use the data and for providing funding for the project. They are also grateful to E. M. de Souza of the Department of Mining Engineering, Queen's University and J. W. Loring of INCO Ltd., principal investigators of the study, for their cooperation. This consulting experience was conducted with the cooperation of Queen's University George L. Edgett STATLAB.

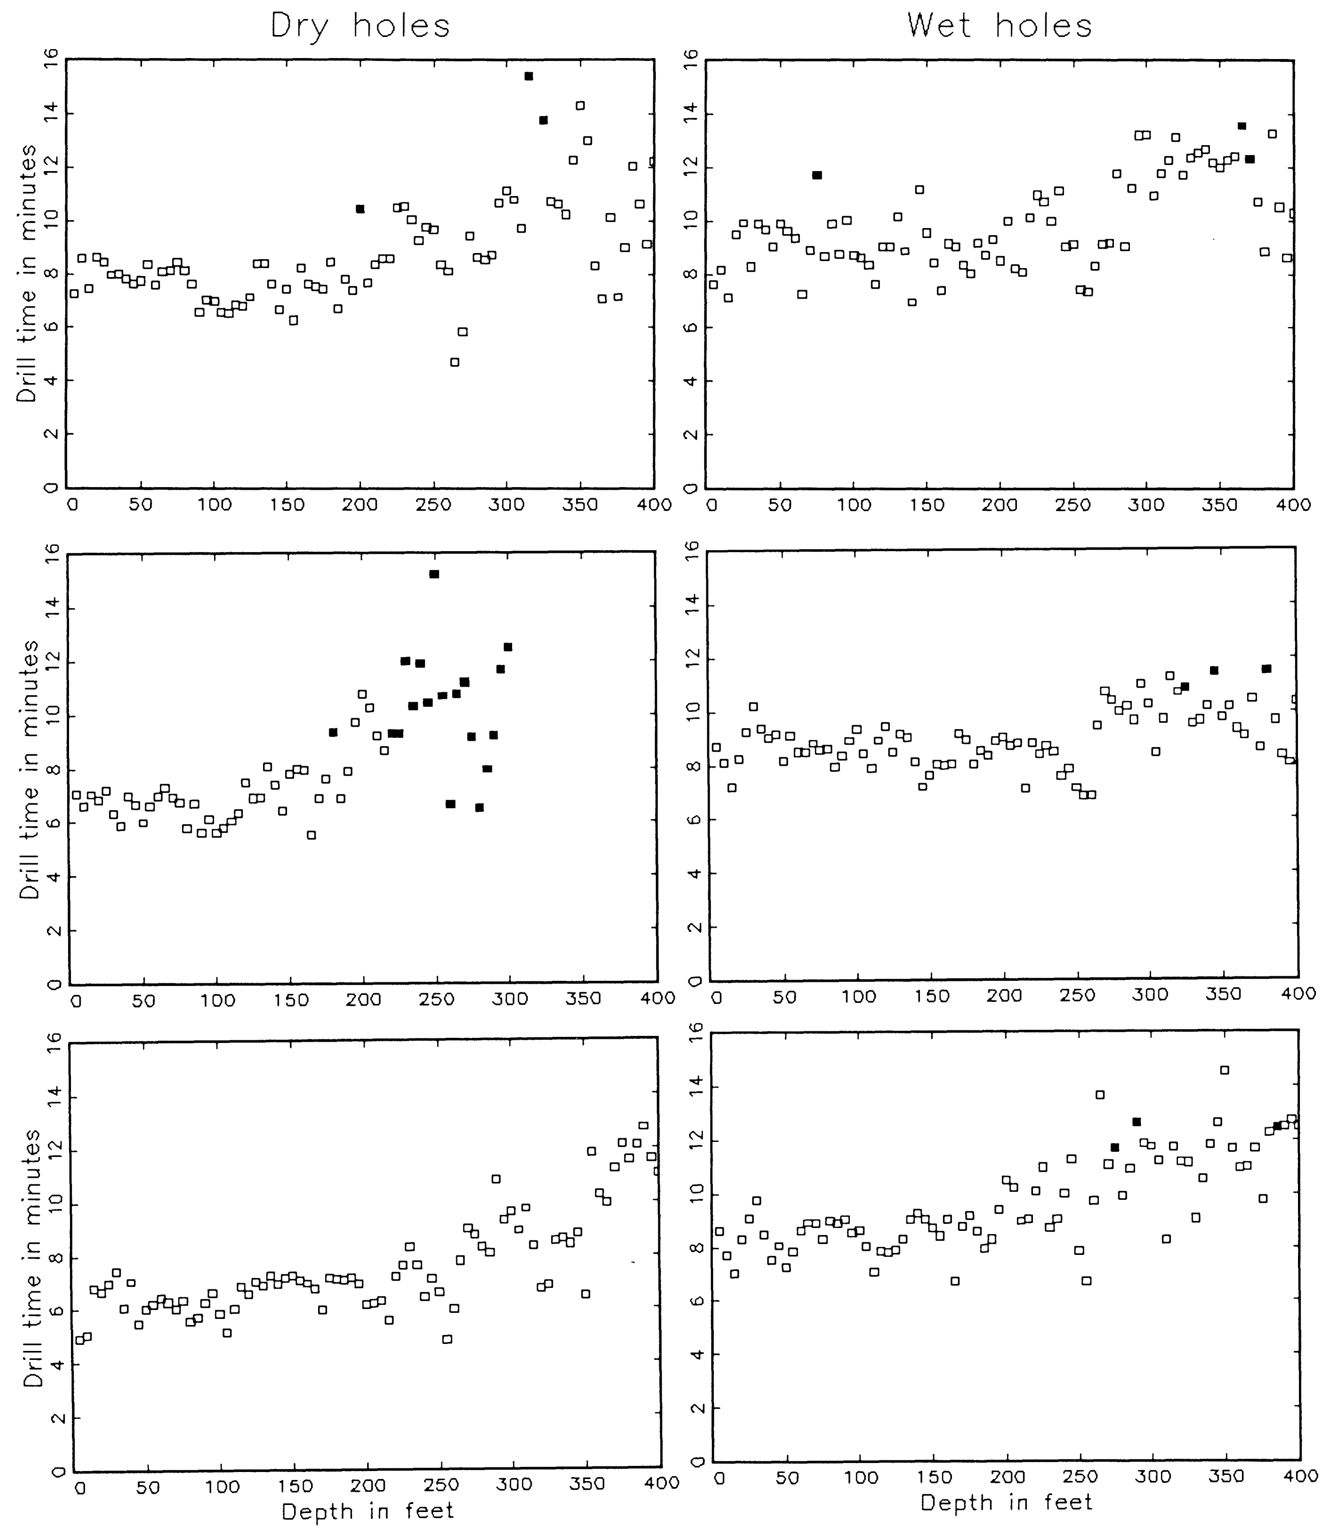

Figure 1. Time (in minutes) to Drill 5 Feet Plotted Versus Depth (in feet) for Dry and Wet Drilling Conditions. Solid squares denote observations for which difficulties were encountered, so the times are known to be too large.

That could have been costly. Why did that happen?

Well, we noticed when we restarted drilling that hole the next day, there was a lot of water which had seeped into the hole. Ground water. And the cuttings just sort of made a kind of paste and clogged up the hole. What we should have done was to blow the hole out for about a quarter of an hour and really dry it out. Then things would have been okay.

Very interesting.

And at about 250 feet we hit an ore "lens."

An ore lens?

Yes, a lens of copper-nickel ore. We could tell because we could really see the drill moving down—normally it barely moves. And the cuttings all started to spew out brown. So, because ore is soft, the drilling times were much shorter.

I can see that. But look! If you line up the plots, you can see that after you hit the ore lens, the times are noticeably larger . . .

Table 1. Drill Times

	Time to drill 5 feet					
Depth	Wet			Dry		
5	7.61	8.68	8.61	7.25	7.07	4.90
10	8.16	8.13	7.71	8.55	6.62	5.07
15	7.11	7.21	7.02	7.43	7.04	6.77
20	9.47	8.26	8.30	8.60	6.82	6.65
25	9.91	9.24	9.07	8.45	7.19	6.99
30	8.27	10.22	9.75	7.95	6.34	7.41
35	9.87	9.39	8.47	7.98	5.88	6.07
40	9.65	9.04	7.50	7.80	6.99	7.04
45	9.02	9.18	8.04	7.62	6.65	5.49
50	9.89	8.14	7.27	7.72	5.99	6.03
55	9.59	9.13	7.85	8.35	6.63	6.19
60	9.34	8.50	8.61	7.58	6.99	6.43
65	7.26	8.49	8.91	8.07	7.28	6.27
70	8.88	8.82	8.89	8.10	6.91	6.03
75	(11.72)	8.55	8.30	8.42	6.74	6.34
80	8.67	8.61	8.96	8.13	5.79	5.57
85	9.87	7.92	8.89	7.62	6.70	5.70
90	8.76	8.33	9.03	6.55	5.60	6.23
95	10.04	8.91	8.53	7.00	6.10	6.60
100	8.69	9.36	8.60	6.98	5.60	5.84
105	8.60	8.44	8.02	6.58	5.77	5.17
110	8.34	7.89	7.06	6.52	6.04	6.03
115	7.60	8.91	7.84	6.82	6.32	6.84
120	9.01	9.43	7.80	6.80	7.47	6.58
125	9.01	8.49	7.88	7.13	6.90	7.03
130	10.15	9.18	8.29	8.37	6.93	6.89
135	8.86	9.01	9.02	8.38	8.07	7.27
140	6.95	8.12	9.24	7.60	7.39	6.92
145	11.17	7.18	9.02	6.65	6.42	7.15
150	9.54	7.62	8.71	7.42	7.80	7.25
155	8.42	8.01	8.41	6.27	7.97	7.05
160	7.40	7.97	9.01	8.20	7.92	6.95
165	9.14	8.02	6.71	7.60	5.53	6.76
170	9.00	9.18	8.75	7.53	6.88	5.97
175	8.34	8.94	9.15	7.43	7.62	7.17
180	8.01	8.02	8.56	8.42	(9.34)	7.11
185	9.15	8.51	7.94	6.68	6.89	7.07
190	8.71	8.35	8.27	7.78	7.90	7.17
195	9.29	8.90	9.38	7.37	9.72	6.91
200	8.50	9.03	10.47	(10.42)	10.76	6.15
205	9.97	8.71	10.19	7.67	10.24	6.19
210	8.19	8.80	8.95	8.32	9.22	6.29
215	8.05	7.10	9.02	8.58	8.66	5.58
220	10.12	8.80	10.06	8.57	(9.28)	7.22
225	10.96	8.39	10.96	10.47	(9.28)	7.62
230	10.72	8.68	8.69	10.53	(12)	8.28
235	10.00	8.46	9.01	10.05	(10.29)	7.59
240	11.12	7.57	10.00	9.25	(11.89)	6.42
245	9.00	7.83	11.25	9.75	(10.45)	7.12
250	9.12	7.13	7.86	9.65	(15.21)	6.62
255	7.42	6.82	6.68	8.32	(10.68)	4.84
260	7.36	6.81	9.70	8.08	(6.67)	5.97
265	8.29	9.44	13.63	4.70	(10.74)	7.77
270	9.13	10.71	11.05	5.82	(11.19)	8.96
275	9.17	10.37	(11.68)	9.42	(9.17)	8.73
280	11.76	9.97	9.90	8.63	(6.52)	8.29
285	9.00	10.16	10.88	8.53	(7.95)	8.05
290	11.21	9.63	(12.63)	8.72	(9.21)	10.79
295	13.19	10.99	11.84	10.67	(11.67)	9.29
300	13.20	10.25	11.73	11.13	(12.49)	9.62
305	10.94	8.44	11.21	10.77	*	8.91
310	11.77	9.68	8.26	9.72	*	9.73
315	12.24	11.24	11.72	(15.40)	*	8.34
320	13.11	10.68	11.18	(18.75)	*	6.75
325	11.70	(10.84)	11.14	(13.78)	*	6.86
330	12.34	9.53	9.05	10.72	*	8.54
335	12.52	9.63	10.54	10.63	*	8.60
340	12.66	10.16	11.81	10.23	*	8.41
345	12.15	(11.44)	12.62	12.27	*	8.81
350	12.00	9.74	14.52	14.30	*	6.46
355	12.24	10.14	11.66	12.97	*	11.79
360	12.39	9.32	10.93	8.30	*	10.24
365	(13.55)	9.09	10.99	7.07	*	9.93
370	(12.28)	10.43	11.65	10.10	*	11.22
375	10.72	8.62	9.77	7.13	*	12.12
380	8.84	(11.46)	12.25	8.98	*	11.53
385	13.24	9.63	(12.43)	12.02	*	12.09
390	10.50	8.32	12.50	10.62	*	12.74
395	8.61	8.04	12.70	9.10	*	11.55
400	10.27	10.35	12.50	12.22	*	11.02

NOTE: Numbers in parentheses are known to be too large.
*Indicates missing values.

Oh yeah! That could be because we went from, say, greenstone to quartzite, and the contact would be a natural place for an ore lens to occur.

And greenstone is softer than quartzite?

Yes.

Well, that's very useful information! So, in your log book, you have annotations of when various things happened, such as when you ran into the ore lens. Do you have information about the type of rock at various types?

No. Unfortunately not. It was easy to tell the ore lens because of the big color change, but not for any of the other depths.

Too bad. Oh well, let's try to see where we stand. The main question is "How much does dry drilling reduce drilling time?" *We have three holes drilled under each condition, with about 80 observations per hole. We know that the drilling times are affected by variations in rock hardness, and so if there were no missing observations, such as occurred with dry hole 2, or where there is extra time due to binding, we could just do a two-way analysis of variance to buffer out that additional variation. Right?*

I guess so.

Good. But we should try to take advantage of all those good annotations and extra information. However, for a start, I'd like you to perform a simple paired-t analysis, and then we'll look at the results of that when I see you next.

Okay.

So what I'd like you to do is to calculate the average for each drilling process at each depth, but just calculate the average for the "good" data values. That is, for the dry holes, calculate the average of the three times at depths where all three times are good, the average of two times at depths where only two times are good, and the single time when only one time is good, and so generate a single column of data consisting of average dry drilling time versus depth. Then do the same for the wet times, and then do a paired-t analysis on those numbers. We'll ignore the fact that some of the averages are more reliable than others for now. Okay?

Yes.

Good. See you next week.

2. AUGUST 28

Hi, Ralph, how's it going?

Fine, thanks. I did the paired-*t* analysis on the "good data" averages and found that dry drilling reduced the time by about 1.4 minutes, which corresponds to about a 15% reduction. And when I converted it to a penetration rate the reduction was even higher, about 19%.

Penetration rate?

Yes. Instead of using the time to drill 5 feet, I converted the values to a drill rate in feet per minute. Actually, I used centimeters per minute.

Well that's very creative! But we'd better be careful about taking reciprocals here—transforming data can sometimes introduce unwanted effects.

Is that right?

Yes. But you have a valid point. We should look at the distribution of the data. We'll come back to that, but first I'd like to try another idea. You told me that you can identify at what depth the ore lens was hit for each hole. Is that correct?

Yes (referring to log book and graphs). For dry hole 1, it was at 265 feet, for dry 2, 260 feet, and for dry 3, 255 feet. For the wet holes, 3 was at 255 feet, and 1 and 2 were at 255 or 260 feet.

So now you're basically using the lowest local drill time near 255 feet, correct?

Yes.

Well, it seems to me that the ore lens would serve as a good index point. That is, maybe we should use the depth of the ore lens to "line up" the data points, rather than the depth from the surface.

What's the advantage?

Well, it's possible that there is a change of rock type after the ore lens, and we can see there was a jump in the drill times there.

Okay. So let's try lining up the data according to the depth of the ore lens. But what data should we analyze—the 5-feet drill times or the penetration rates?

We'll stick with the drill times for now and have a look at the residuals. While you do that, I'll think about which metric we should be using, and how we should attend to missing values. Actually most aren't missing, we just know they're high.

3. AUGUST 29

Hello, Dr. Watts.

Hi, Ralph. Well, I analyzed your data with a view to deciding whether we should use the times for 5 feet or the penetration rates. What I did was similar to what I asked you to do. That is, for both wet and dry holes I calculated the averages of good data points and then calculated the deviations of the good data points from the averages. Then I did a normal probability plot of the deviations. Then I did the same thing but using the inverses of the good data points. It turns out that there is really no reason to choose one metric over the other—all the normal probability plots were extremely straight. In fact, if anything, the penetration rates were more normal!

Good. If it's okay, then, I'll use them.

Fine. I also tried using the ore lens to line up the data for the paired-t analysis. When you did the paired-t analysis, you got a difference in the times to drill 5 feet of 1.396 minutes with a standard error of .125 minutes. By lagging dry hole 2 by one rod length and dry hole 3 by two rod lengths, and lagging the wet holes by two rod lengths, my analysis gave a mean difference of 1.374 minutes and a standard error of .120 minutes, so the alignment did seem to provide a more sensitive analysis.

That's interesting.

Thanks. For the penetration rates, in cm per minute, the no-lag difference (dry–wet) was 3.35 cm per minute with a standard error of .29 cm per minute, and for the lagged penetration rates, the difference was 3.29 cm per minute with a standard error of .27 cm per minute, so again the standard error decreased.

Good.

And for both analyses, we get a nice tight confidence interval well away from 0—that is, drilling dry does speed things up. Now tell me, Ralph, why would dry drilling be faster?

Well, I think what happens is this: the bit is busy breaking off cuttings, and the high-pressure fluid is supposed to carry the cuttings up the hole. When air is used as the fluid, it only has to lift the cuttings, but when water is used the fluid has to lift the cuttings *and* the water. So the cuttings can't be lifted as easily with the water added as with the air alone. And if a large chip breaks off, then that can't be expelled as easily with water, so the chip either impedes the flow of cuttings up the hole, or falls back down and gets hammered by the bit again. In either event, the drill is doing more unproductive work, regrinding cuttings.

Makes sense. It's always nice to be able to give a physical justification for a result. So, now we have to worry about missing values and unreliable readings. I did do some work, and basically I just made your paired-t analysis a bit more efficient by using weighted least squares. I took differences of the averages, as you did, but I took into account that each difference has a variance proportional to the sum of the reciprocals of the number of good values in each average.

And what happened?

About what you would expect since there weren't too many missing observations. Except for dry hole 2 at depths greater than 200 feet, the results were very similar to the crude paired-t analysis. Using the data as is, that is, not using the ore lens to align the data, the difference (wet–dry) for 5-foot drill times was 1.391 minutes with a standard error of .115 minutes, and for the penetration rates, the difference (dry–wet) was 3.41 cm per minute with a standard error of .28 cm per minute. Aligning the data according to the depth of the ore lens gave results of 1.403 minutes with a standard error of .115 minutes for the 5-foot drill times, and 3.43 cm per minute with a standard error of .26 cm per minute for the penetration rates.

So there's not a helluva lot of difference in the results using the weighting.

Correct. But the aligning and weighting have produced a more sensitive analysis.

That's good. But there is still one other thing . . .

What's that?

I'm still interested in the fact that the penetration rates tend to approach one another as the depth increases, so the advantage of dry drilling may be decreasing.

Hmmm . . . Is there any reason for that?

Yes, it's reasonable. As the hole gets deeper, the cuttings have to be expelled from a greater depth, so it is harder to lift them. Then more pieces would fall back and be reground, so the efficiency would decrease . . .

And the wet drilling would not be so affected by increasing depth?

Well, I don't think that it's not affected, I just feel that it is possible that the difference in flushing ability isn't as great when the depth increases. It is more likely, though, that the convergence is due more to dry drilling difficulties that were encountered.

So there could be a depth effect. And you'd like to account for it.

Yes. I would like to at least explain, mathematically, the apparent convergence. Any physical reasons will have to be developed later.

Hmmm. Well, in our previous analysis, basically all we concerned ourselves with was finding out how much greater the dry hole penetration rate was compared to the wet. So we looked at the differences and didn't really pay attention to the averages.

Right. But we could see, from the plots, that there was a slope to the data.

Well, to the dry data, yes, but the wet data was more like a horizontal line with a jump after the lens. However, we could confirm this by fitting a model to the wet data as follows: Let x denote depth in cm, and x_0 denote the depth of the lens, also in cm, and let's introduce an indicator factor, say x_1, which is 0 for $x < x_0$ and 1 for $x \geq x_0$. Then the penetration rate in cm/min is modeled as

$$w = \beta_1 + \beta_2 x + \beta_3 x_1 + \beta_4 (x - x_0) x_1,$$

where β_1 represents the nominal wet drill penetration rate (cm/min), β_2 gives the change in penetration rate as the depth increases, or the slope, (cm/(min cm) = 1/min), β_3 gives the "step" change in the rate below the lens (cm/min), and β_4 gives the change in the slope below the lens (1/min). We could fit a similar model to the dry drill data, but it's better to combine all the data together, by introducing yet another indicator factor, say x_2, which is 0 for wet data and 1 for dry. We stack the wet data on top of the dry data and call that y and fit the model

$$y = \beta_1 + \beta_2 x + \beta_3 x_1 + \beta_4 (x - x_0) x_1 + x_2 \{\delta_1 + \delta_2 x + \delta_3 x_1 + \delta_4 (x - x_0) x_1\}.$$

What's the advantage of that?

Well, it uses all the data at once, which makes it efficient. And all the parameters which occur with x_1 and x_2 are change *parameters—for example, β_3 has the interpretation of the change in penetration rate with wet drilling because you are below the lens, β_4 is the change in slope with wet drilling below the lens, δ_1 is the change in penetration rate due to dry drilling, δ_2 is the change in slope due to dry drilling, δ_3 is the change in penetration rate below the lens due to dry drilling, and δ_4 is the change in the change in the slope below the lens due to dry drilling.*

I see.

But, and this is very important, we automatically get the standard errors associated with each of these change parameters, and so we can easily see whether a change parameter could be 0—we don't need to calculate differences between parameters and find out the standard error of the difference based on the standard errors of the individual parameters, and all that stuff. And of course, we'll want to do weighted least squares on the averages. And account for the lens data too.

How will you do that?

I think I'll just edit out the few values corresponding to the lens data in each hole rather than introduce some more parameters to account for those anomalous values. And I'll use the lens depth to align the data as well.

Pretty involved analysis . . .

Yes. It's amazing how a simple problem can escalate into a complicated solution. Anyway, I'll bash these numbers out and we'll see how they look.

Great.

4. SEPTEMBER 8

Hi, Ralph.

Hi, Dr. Watts. How're the results?

Very interesting, I think. I hope you'll be pleased.

I'm sure I will.

The regression coefficients and their standard errors for this first model (1) are shown in the table (see Table 2 on page 9). From this, we can see that both the slope (β_2) and the slope change (β_4) for wet penetration rates are likely 0 since they have t ratios of -1.4 and .05.

So the regression analysis confirms our suspicion that there was no linear change of wet penetration rate with depth?

Right. But we're told more. The coefficients for the step change due to dry drilling (δ_3) and for the change in the change in slope due to dry drilling below the lens (δ_4) could be 0 since they only have t ratios of 1.05 and $-.49$.

So we could drop four terms from the equation?

Yes, but I thought we should proceed carefully, so I only removed the slope and slope change for the wet penetration rate. The regression results changed to those for model 2. Now there is only one term which could be dropped—the dry slope change below the lens (δ_4)—since it has a nonsignificant t ratio. There were also a half-dozen largish standardized residuals, including one biggie with a t ratio of 3.7, so I removed that one observation and refitted the model excluding the dry slope change factor, to give the results shown as model 3.

Not much difference there.

No. One thing which surprised me a bit is that a plot of the residuals versus depth didn't reveal any noticeable change in variance below the lens. From the plots of the data, I thought that there might be.

So, how should I present the results?

Well, I suggest you give them the final equation with the parameter estimates and their standard errors, and then set out a little table which shows the improved penetration rates at, say, every 100 feet to a depth of 400. For example, at 100 feet, the dry penetration rate is 21.9 cm/min and the

Table 2. Parameter Estimates for Regression Analysis

	Model 1		Model 2		Model 3	
Parameter	Estimate	Standard error	Estimate	Standard error	Estimate	Standard error
β_1 (cm/min)	18.1	.38	17.7	.19	17.7	.19
β_2 ($\times 10^{-4}$/min)	−1.26	.89				
β_3 (cm/min)	−2.70	.72	−3.42	.36	−3.42	.35
β_4 ($\times 10^{-4}$/min)	.14	2.82				
δ_1 (cm/min)	5.54	.54	5.99	.43	6.10	.40
δ_2 ($\times 10^{-4}$/min)	−4.53	1.28	−5.78	.89	−6.10	.82
δ_3 (cm/min)	1.07	1.02	1.78	.81	1.39	.69
δ_4 ($\times 10^{-4}$/min)	−1.97	3.94	−1.81	2.82		

wet penetration rate is 17.7, giving a 24% improvement. That's pretty good!

Yes, it is. In fact, I did some calculations for different improvements, and for a mine working 3,000 tons per day, with this kind of improvement the savings in labor and parts would be about 23%, which amounts to about $90,000 per year.

Well, that's certainly a worthwhile result. Congratulations.

Thanks.

There is one minor point, which I should bring to your attention . . .

What's that?

Well, the residuals are pretty well behaved, but they do show autocorrelation.

And what does that mean?

Well, in your case I think it means that the rock you were drilling through was not completely homogeneous, so there were sections where the drill moved more slowly, and others more quickly, so the drill times kind of meandered around the nominal rate, rather than being completely randomly scattered around it. We could take account of this, but the main effect would be to reduce the parameter standard errors a touch, without affecting the conclusions, so I think we'll just call a halt to the analysis here. Okay?

Suits me.

Good. By the way, what about the drill-life data?

Oh yes! Well, what we did was measure how much the buttons were worn down after a depth of 400 feet. For the dry holes, the wear was 16 and 18, and for the wet, 60, 55, and 50. In 25-hundredths of an inch—not much eh?

No, it sure isn't.

But it makes a big difference economically as the bits will last longer and will drill more holes. Also, since they don't have to be sharpened as often, you can drill without pulling the rods as many times.

I see. Dry drilling not only speeds up the actual drilling procedure, but can save time because the bit doesn't need to be resharpened as frequently.

Right. In fact, a good deal of the $90,000 is due to increased bit life and savings in sharpening time. There are also other advantages, such as a cleaner work site, reduced maintenance costs on equipment, less cleaning and easier loading of blast-holes, and less slime accumulation in the ditches, and even reduced ditch cleaning.

Gee, isn't it nice that you get so many advantages.

It sure is. Anyway, I'd better get going. Thanks again for your help.

Thank you. I've enjoyed working with you.

5. PROLOGUE

As with most real problems, there is no one unique correct solution. Depending on one's tastes, the problem could be treated as a problem in multivariate data analysis, in time series analysis, and indeed, in multiple time series analysis, for example.

We hope this article will spawn more articles of the same type. It would be very interesting to be able to read about other problems with the full background—the physical layout, the practical and/or economic motivation, and so on—so that we could adopt the problems and data sets and use them for motivating students in our own classes. To receive a copy of the data by electronic mail, send a request to WATTSDG@QUCDN.Queens U. CA

[Received October 1989. Revised March 1990.]