

Telephone calls and the Brontosaurus

Adam Atkinson (ghira@mistral.co.uk)

This article provides more detail than my talk at G4G11 with the same title.


I am occasionally asked questions along the lines of “When do you ever use any of this stuff in real life?” or “What is the hardest mathematics you have ever used in real life?”. I imagine other G4G people have had similar experiences.

An answer I found myself giving to both of these kinds of questions recently also has the moderate advantage that it uses material which is, or could reasonably be regarded as being, “secondary school maths”, so it might be of some interest or use for people whose “this stuff” is high school level rather than university level maths.

I was working for a company which installed/maintained internal telephone systems for organizations of various sizes, including the links between these systems and the outside world. Note that what I did at this company did not involve telephone systems so during the events of this story the whole situation was new to me. What I did may not necessarily reflect best practice on the part of people who really do this sort of thing for a living, but as with spherical cattle or drunks and street lights, we might be willing to sacrifice some accuracy/plausibility for the sake of creating a more accessible exercise. Also, and principally, I don't want hundreds of people to think that anything I say here reflects that company's actual approach to a problem like this.

I was approached by a manager and told that one of our customers felt that we were charging too much for phone calls. The costs worked like this: incoming calls and internal calls were free. Only calls to the outside world cost money so from now on we shall only be considering outgoing calls. And we shall only consider standard outgoing calls (to national numbers, not international ones for example). A call cost (some fixed amount of money) as soon as it was connected, and if it lasted at most (some duration) there was no extra charge. If the call lasted longer than that, there was an additional charge of (some amount of money) per unit of time.

For the sake of argument, let's say that our price was 5p for any call up to 3 minutes, then 1p a minute after that.




The customer had 5 quotes from our competitors, all expressed similarly. A for the first B minutes, then C per minute after that. The customer felt that our tariff was “clearly” too expensive. It was not “clear” to us why it was, so we asked, and were told that our value of C was too high. I was asked by my side if this was reasonable, and how I would compare such tariffs. As it happens, our value of A was smaller than some of the others, so I asked if the customer made enough “long” calls for our high value of C to cause a problem. It turns out neither side really knew, but since we ran the phone system it was of course perfectly possible to get a log of all calls going back months to look at this kind of thing.

With a list of call lengths, we can compare tariffs by seeing what each tariff would expect one to pay for that set of calls. However, one might not have such a log, or it might be so short that it might not be considered to be representative. Or the customer might fear that over time the length of calls might change.

Can we make any attempt at all to compare tariffs without a huge log of calls? I am asking this rhetorically, so clearly the answer must be “yes”.

For starters, if C is too large then the cheapest way to make a very long call would be to hang up and re-dial every B seconds. This would be an annoying thing to have to do but one could imagine some people going to this much trouble. Certainly if modems were still a thing and costs were like this I would expect people to arrange for their modems to behave in this manner. Let's assume that even if some of our tariffs are like this, real people are not going to bother to redial all the time and will pay the tariff rate.

Any easy comparison is one that looks like this:




Clearly as the customer we would choose the yellow tariff here.


Let's assume any tariff which, like the red one here, is totally undercut by some other tariff is removed from further consideration.

Two other things can happen although the difference between them probably doesn't matter much.

In the first case one tariff could be cheaper for short calls and the other for longer calls:




In the second case, one tariff could be better for medium length calls only.



In these cases, you need to know how many calls of each length the customer makes.

Whilst in principle the probability distribution of call lengths could look like almost anything, perhaps in real life it can be treated as coming from some family with a small number of parameters. As it happens, I knew what family of distributions was commonly used for modelling purposes at least in some exercises about such things, but I didn't know how real-life it was. So let's not mention this again until later.

One might suppose that the probability distribution of phone calls might be brontosaurus-like. As A. Elk put it, the brontosaurus is thin at one end, much much thicker in the middle, and thin again at the far end:




Or perhaps more of a stegosaurus. I think we can all agree that telephone calls can't have negative length, but it might seem plausible that there's some common medium length and calls longer and shorter than that are less common.

If we know what this graph looks like then we can calculate the average cost of a call as:
 $A + C * \text{Prob}(\text{Call lasts at least } B) * (\text{Mean additional duration of calls which last at least } B)$,
 since all calls get charged A immediately, then some get charged more.

If we have a nice formula for our probability distribution we can turn this into something with integrals in it.

Calculating the mean additional duration is going to be possibly quite annoying. Wouldn't it be nice if it weren't annoying? The “spherical cows” assumption at this point is that the distribution of call lengths is exponential, because then the average additional duration of calls of length at least B is the same as the mean of the distribution as a whole. Since I have information about hundreds of thousands of calls made in real life, though, let's look at that:



Which is nothing like a brontosaurus at all. The distribution sometimes used in exercises about this kind of thing is the exponential, which has only one parameter, so if you know the mean you know everything you need to. Of course, to find out a customer's mean call length you need some call logs and if you have those you could run the calculations based on those as mentioned earlier.

Now actually this graph is taken part way through cleaning up the call logs, and it seems possible that some information about the “cleaning” might be of interest.

A similar histogram made from the raw logs had a peak at 1 second rather than 0 seconds, which would have ruined the “it's exponential!” impression. I thought maybe this was because calls from 0.5 to 1.5 seconds were being called 1 second calls, so perhaps the 0 second calls had a very narrow time range. Actually it was stranger than that. In the call logs, if a call started in one calendar second and finished in the next one, it was called a one second call, even if it was actually, say, 0.1 seconds long, and this handling of “calendar seconds” clearly pushed many calls into seeming a substantial fraction of a second longer than they really were. Of course, for most purposes an error of under 1 second wouldn't matter, but we're worrying pointlessly about the relative heights of 0 and 1 second columns on a histogram here so let's try to fix this.

Fortunately, the call logs in the raw data used for this graph also contained information from which the length of the call in 50ths of a second could be deduced, and it is using that rather than the “duration” column that the above histogram was produced. Incomplete seconds are rounded down, so 0 to 49 50ths of a second count as 0 seconds, etc.

Unfortunately, the 0 second call is now incredibly tall. Can we find an excuse for making it shorter again somehow? Well, yes. In the call logs from the original story I was given a log only of outgoing calls but I used a log of all calls to make this graph. I ought to remove internal and incoming calls from it, and one particular class of incoming call that shows up as being 0 seconds long is a call which is diverted automatically to voicemail. For some reason, such calls show up as a 0 second call to the phone followed by a real call to the voicemail system. Since these calls are incoming, they should be eliminated along with internal calls, international calls, calls to freephone numbers and so on.

Of course, I have not shown that the graph above really is exponential, merely that it looks closer to an exponential than to a brontosaurus. Using a log scale on the y axis would be the sensible thing to do here:

