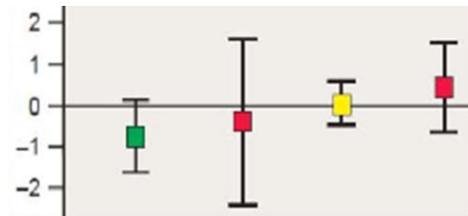


A truly smart predictive dialler!

What are the chances?

“Well actually the chances are 100% with a Standard Error of 0.0!” – Tom Thompson – Senior Software Architect at InVADE

Statistics are the bread and butter of Market Research, but the basic numbers that your surveys churn out are, on their own, not enough. They need to be qualified. You need to know how accurate those statistics are; that's why surveys have quotas, to ensure there is enough sample to validate them. For Market Research statistical results are unusable without knowing their accuracy, and the skill of the Market Researcher is in knowing how accurate each one needs to be.



For diallers it's the same. At first glance predictive dialling algorithms based on churning call and agent statistics into probabilities may appear to be mathematically rigid and over simplistic. However, calculating those probabilities and running them through some “what-ifs” is the easy part. For each of the primary statistical values that the dialler is continuously recalculating there are a set secondary, companion values no less important. These denote **accuracy**, **fluctuations** and **patterns**. It is these values that determine if, when and how the primary values are applied. At the core of the smart predictive dialler are subtle judgments based on these secondary statistics.

What I want to show in this white paper is three things. Firstly, that some of the most fundamental performance factors of a dialler are qualitative, not quantitative. Secondly, that you can still use a mathematical approach to make sense of these as long as you clearly understand how these qualitative factors affect the dialler's stats. Lastly you can intelligently apply some of these secondary statistical values to get the best out of a dialler.

“Statistics are like bikinis. What they reveal is suggestive, but what they conceal is vital” - Prof Aaron Levenstein

Predictive Dialling

Key Variables

"The mean number of legs per person is a little less than two, and most of us have exactly two, so most of us have more legs than average" - The Open University, UK

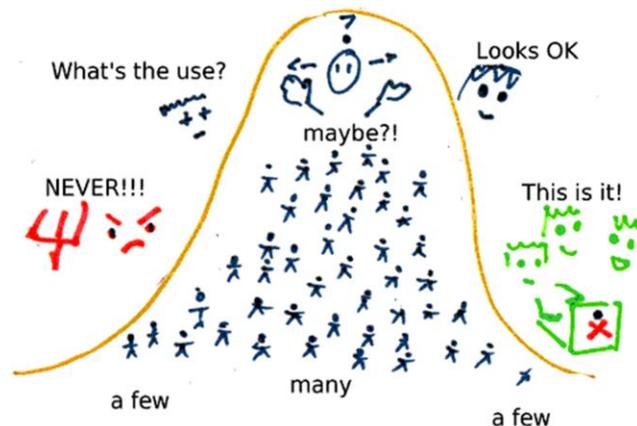
Dialling calls from a sample set for a group of interviewers produces lots and lots of stats; some are direct indicators of key factors and others are just performance outcomes. It's not always easy to understand which factors and values identify 'causes' and which 'effects', and the relationship between them can sometimes be complex and counter intuitive.

The discussion below is on some of the most fundamental factors.

List quality: At the heart of a dialler's performance parameters is the connection rate (% of calls that connect to respondents). The connection rate is the result of number of factors but most fundamental to it is the quality of the call sample list. At one extreme you might have random digit dialling or a remainder list (built of calls that "rang-out" or "busied-out" multiple times). You would expect to get very low connection rates from such lists. At the other end of the scale you might have lists of scheduled or appointment call backs such as those in panel based projects.

However connection rate is not the limit of the list's effect on a dialler. Just as important can be the type of call failures that a list produces. This is especially so when the connection rate is low. In cases where the bulk of calls are failures the performance of that predictively dialled project will depend heavily on how quickly the calls fail. For instance compare a project of 10% connections and 90% bad numbers against a project of 20% connections and 80% no answers. Even though one project has just half the connect rate, it could produce shorter interviewer wait times and more talk time. This is because for such project the calls that fail do so in just a few seconds, as opposed to calls that ring and don't answer and so clog up your trunks for around 10 times longer. This is why random digit dialling projects will tend to keep interviewers busier than projects running on remainder sample.

Number of Interviewers: On the face of it, the number of interviewers on a project, or more precisely, the number of free interviewers (ready to take a new call) is a simple quantity and so should plug neatly into the diallers calculating engine. The mathematics works like so: the more free interviewers you have the bigger the area you get under the Bell Curve where the number of connected calls does not exceed the of free agents. It basically means that the more free interviewers you have the more the over-dialling (for each and every free interviewer) you can do without increasing the risk of producing



Predictive Dialling

abandoned/silent calls. So the dialler may calculate that when you have 2 free interviewers can have just 4 calls progressing (made but not yet connected/failed), but with 4 free interviewers you can have 12 progressing calls for the same risk of producing a call that must be abandoned. In the first case the over-dial rate is x2 and in the second its x3. This is why predictive dialler manufacturers will regularly and rightly tell you that you will see greater productivity gains with more simultaneous interviewers on to a project; much better to have 40 interviewers working for one hour than 10 for 4 hours.

That's the simple approach to interviewer numbers. From there you can start to get smarter. You can bump up the number of free interviewers with the number of "potentially" free interviewers: that is the number of interviewers that are likely to become ready to take a call in the near future. This is perhaps the area of greatest subtly in a smart predictive dialler where data accuracy, fluctuations, and pattern matching are applied to the greatest extent.

Lastly interviewer numbers can have an indirect on the application of other dialling factors. A high number of interviewers can be used to offset low data accuracy elsewhere. Such situations can produce 'turbo-charge' effects where the dialler can determine that it can make decisions based on the overall accuracy of all its factors rather than limiting it by the weakest piece of data it has.

Survey: The length and detail of a survey determines how long interviews can spend in calls. Of course not all connections will run an interview from beginning to end but projects with long interviews will increase interviewers "talk-time" (time connected to respondents).



He's Pattern Matching

What's the effect of that? Well, interviewer talk-time will be a higher percentage of interviewer activity and, in consequence, waiting time will be a lower percentage. But what's the effect on the dialler's performance? The more time interviewers are spending in calls the less time they are spending waiting for calls, so at any one time the average number of 'free' (or 'waiting') interviewers will be reduced. This will throttle back the predictive dialler's desire to over-dial. So even though overall waiting times may be reduce, an interviewers' wait for each call will be longer. This produces a counter intuitive perception that the interviewers are actually waiting more. They are not. They are waiting longer, but far less often.

It is in the nature of many MR CATI projects to produce in many short calls (refusals) and a relatively small proportion of long calls (complete interviews). Overall an interviewer may spend the majority of *talk-time* in good interviews, but these 'good' calls represent a minority of the connections they get.

Predictive Dialling

Call time distributions will form a classic “double peak”. Recognising *patterns* such as these can, in certain circumstances, significantly improve dialler performance.

Telephony Network: As discussed in “List Quality” the distribution of call failure types can have a huge effect on a project. While this distribution may be largely a consequence of the list quality, it can also be skewed by the quality of information available over the Telco connection.

I won’t go into detail regarding all of the problems that can arise from telephony network limitations – that could comprise a white paper of its own – but determining call failure reasons can be fraught with issues. It can depend on where in the world you are dialling from, where you are dialling to, and who you are getting your public telephony connection from.

Solving such problems can be a fine judgement, and on occasions they can be more expediently solved by using good approximations, rather than trying to achieve complete accuracy in digital signalling and in-band processing (what’s on the line). Sometimes complete accuracy just isn’t possible.

Reflecting on factors like the four above you will have come to realise smart predictive dialling needs to consider more than just clean and clinical numbers. Smart dialling is not simply about dialling sample, getting a connection rate, and re-processing that rate to calculate how many calls should now be made. Some of the dynamics of the dialler environment are more ‘fuzzy’ than others, and need a more circumspect approach. Others, like the number of interviewers on a project have straight forward influences, but may also have subtler ones, and can be used in conjunction with less quantitative data to make judgements. Lastly, there are connections and feedbacks, positive and negative, which need to be understood and communicated to the Market Research CATI Manager.

Making Judgments

Accuracy

What do I mean by accuracy? The accuracy of a statistical result can be considered as the inverse of the well know statistical term “Standard Error”. Standard Error of the mean (SE), as many of you already know, is related to the distribution of values and the sample size. Hard core stats goes way deeper and starts using Greek letters pretty quickly, but this is as deep as I’m going. The bottom line is: there is a way of knowing how inaccurate your stats are, and by corollary, how accurate they are.

$$SE_{\bar{x}} = \frac{s}{\sqrt{n}}$$

Predictive Dialling

Predictive dialling collects and uses all kinds of statistical values, for example: connection rate, various failure type rates; average time to answer; average call length; average time in wrap-up (between ending one call and requesting a new call). Knowing the SE allows the dialler is to set thresholds on when stats are accurate enough to be used. Just as importantly, SE can be used to set sensible limits on optimistic and pessimistic calculations. This later mechanism can be very powerful when targeting to or limiting by an abandoned call rate. How close a dialler is to the target or limit determines how much room there is “to play with”. That’s all well and good but then you need a sensible mechanism to act on that. Using the SE gives you that.

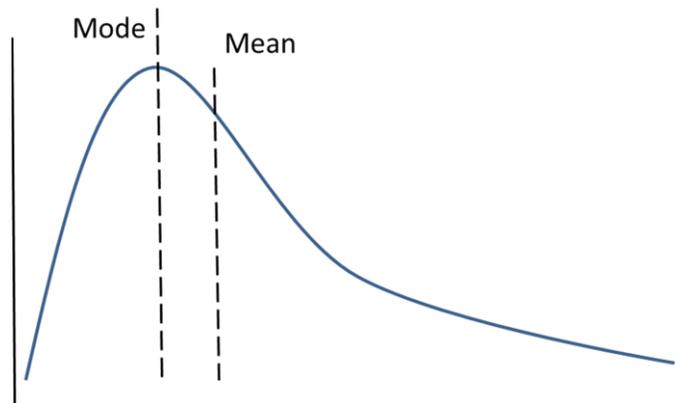
Fluctuations

By the term fluctuations we are really covering the notion of range and distribution of values. This often gets condensed in to the statistical term “Standard Deviation”. One of the ways a dialler can use the notion of fluctuations is in identifying what is currently within normal boundaries and what is not.

Take for instance task of calculating a count for “potential Interviewers” that was introduced earlier. Part of this calculation will determine what proportion of interviewers currently in wrap-up state might be free very soon. This is easy; let’s work out the probability and times it by the number of interviewers in wrap-up. Pop, there’s your number...Err no. There is an interviewer who has been in wrap-up state for way longer than normal. There will likely be a good reason for this, maybe they have an issue that needs supervisor discussion; regardless, the dialler must discount this interviewer.

Patterns

The previous section describes the use of identifying normal behaviour and beyond normal behaviour. On the one hand we can in real time discount any interviewer (or call) showing behaviour that deviates too far from the norm; but that’s not enough. If a value is determined as being not normal, it should also be excluded from the (calculation of the) average that is the

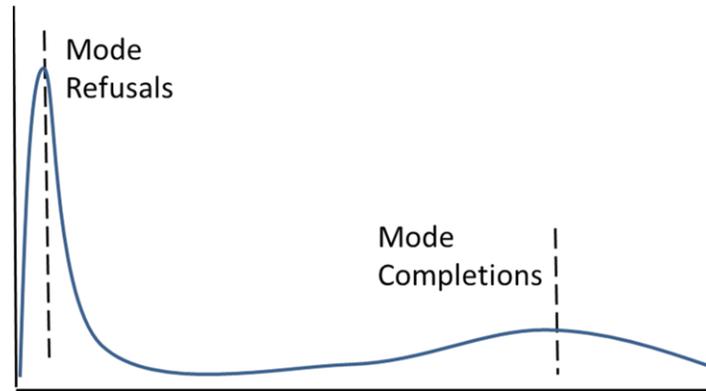


basis for what we are defining as normal. This is the jumping off point for identifying what type of distribution we have and if it fits a well-known pattern for the distribution.

The example above would display a distribution curve where the mode of the distribution (peak) is significantly less than the mean or average value (shown above here). However there are other patterns very common to MR CATI and just as useful to identify. One Classic example is the call length distribution that produces a double peak. In projects where conversion rates are low we usually see a high number of short calls (refusals) and a smaller number of longer calls (completed interviews). The double hump (shown below) denotes two clearly distinct call types that can be treated as such.

Predictive Dialling

Within the invade dialler many of the key stats are submitted to distribution pattern recognition and complementary treatment. The sum of all of this is to squeeze an extra few percentages of performance out of routine campaigns, but just as importantly it can separate, not just individual events, but whole categories of abnormal behaviours. This, for example, allows the dialler to perform predictively, even in mixed mode dialling (preview and group dialling in the same campaign).



Gears

There is no such thing as 0.48 of a connection. There is no such thing as 0.83 of a waiting agent. When the dialler is predicting a range possible call outcomes, there is no continuous spectrum. Calls will connect or not connect and nothing in between. The result is a known as a binomial distribution. The effect on the dialler is perceived as “step change” just like the gears in your car.

However there is another type of gear know in engineering as “Constant Torque” through which power changes are applied smoothly and continuously. The Invade dialler is able to do the same on two levels. As average values are established and change by progressively smaller and smaller amounts, the dialler is able to track these changes. You would expect nothing less. However the dialler is also able to do the same in respect of the accuracy (Standard Error) of these calculated values and gradually increase their influence as their accuracy increases. This ensures the dialler ramps up its performance smoothly from a standing start, and does not overreact to large changes early in its life.

Conclusion

This White Paper does not need to teach Market researcher the power of statistics. But statistics can be used as brutally as a sledge hammer or as delicately as like a painter’s pallet. I hope in this discussion you get an idea of the approach Invade has taken to constructing its dialler. And this approach is applied again and again to all aspects of the dialler. The examples given here just scratch the surface. The Invade Dialler is not just a calculating machine. It’s a decision making engine that uses many levels of statistics both subtly and smartly to give the best performance possible.

“A statistical analysis, properly conducted, is a delicate dissection of uncertainties, a surgery of suppositions.” – M.J.Moroney