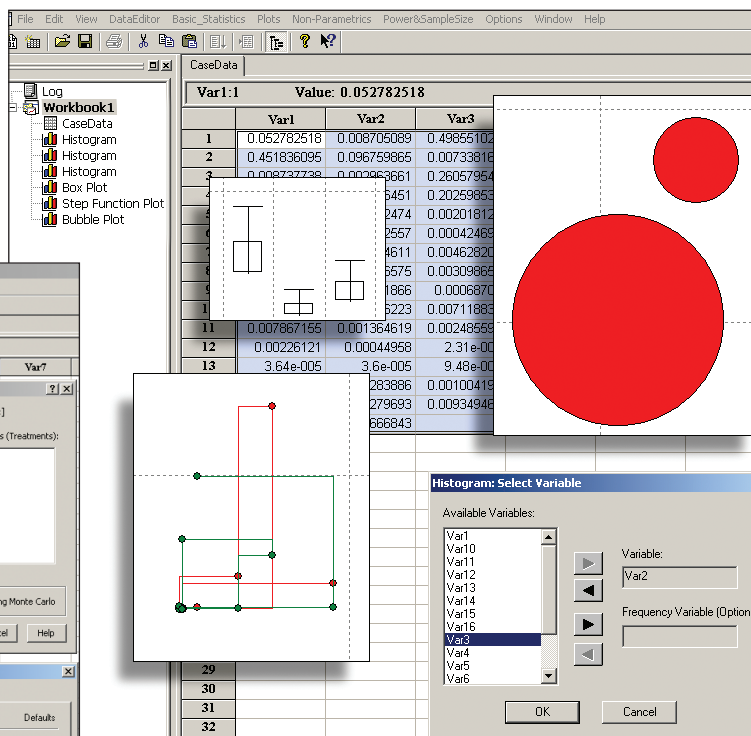
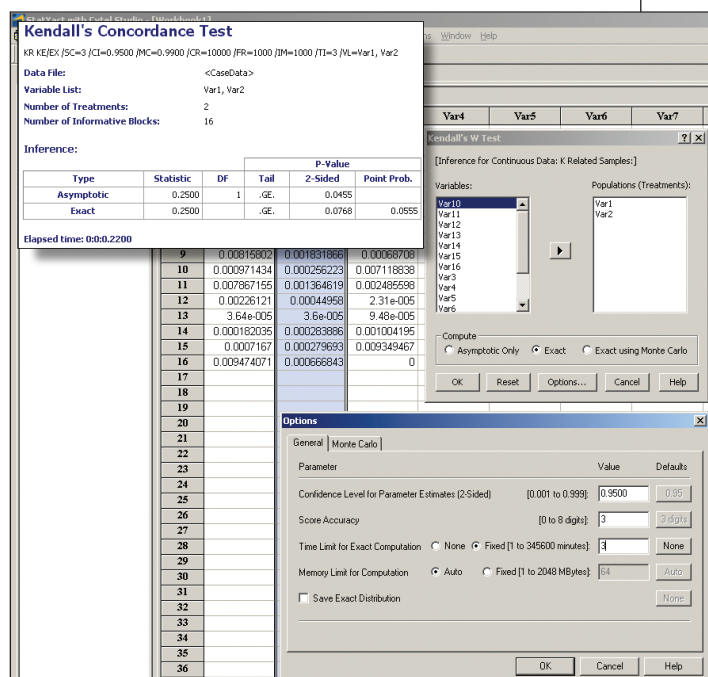


Non-parametric statistics without tears

Exact measures in statistics were painful measures, when *Felix Grant* was a student. Not any more, thanks to StatXact 6 with Cytel Studio



Long ago, when I was a student, when a computer and its high priests occupied their own air-conditioned building far from the reach of undergraduates, we had two eccentrically flamboyant statistics lecturers.

Professor M, with whom we navigated the theoretical mysteries of classical distributions, refused to use the words 'non-parametric statistics'. If forced to acknowledge the subject at all, he would spit out 'quick and dirty statistics'. If you can't establish the underlying distribution, he would say, then don't do the statistics.

The task of guiding us through non-parametric methods fell to Dr P. In contrast to her colleague's reactionary venom, her enthusiasm bordered on sadism. Many a litre of midnight oil was burned, as we laboriously enumerated every possible permutation from a smallish results table

Thirty years in the real world, and the arrival of cheap, readily available computing power, have pushed both of these characters into the rosily humorous glow of history. From time to time, though, they are called back to mind by some discussion or other; and exact measures is one such topic. With the benefit of hindsight, from a safe distance, I can see where both were coming from. A feeling for the relationship between theoretical and exact distribution results is what Dr P was trying to instil in us – and was also the answer to Professor M's unbending scepticism. I only wish that I had been able to see it at the time.

Specific exact tests are, of course, around for particular purposes. They appear in generic statistics packages, though Monte Carlo simulations of known theoretical distributions

with given marginals, before she would tell us which one we were to test against a theoretical distribution.

(courtesy of all that cheap and readily available computing power) have reduced concerns about assumptional errors. Individual exact analyses have long been programmed for particular purposes in particular contexts. StatXact 6 with Cytel Studio (hereafter StatXact), however, is the first open-market product at which I've looked that makes exact, non-parametric measures its raison d'être.

The philosophy of a statistical analysis product can be read in a glance from its documentation, whether paper or electronic. The emphasis may be on a painless leap from data to result, or on making clear to the user exactly what is happening and why. StatXact is certainly easy and painless to use, but the documentation is of the second variety – a model of clarity and integrity that I wish had been available to Dr P all those years ago. I tried it on a mathphobic acquaintance, who grasped the ideas involved immediately.

Transparency and accessibility are a vital component in the acceptance of any product in these GUI-acclimatised times. The publish-

ers are obviously aware of this because they make a major feature of the interface – to which that tag in the product name ‘with Cytel Studio’ is a reference. Studio is an environment in which most users will feel at home, these days – menu and flat tool bar at the top, a narrow vertical pane at screen left, holding a project tree from which material is loaded into the main work area. A case data grid occupies the main area at start up.

Apart from the usual Windows options, such as file, view, and so on, the menu offers a data editor and four specific stats headings. Non-parametrics are, at core, really what the package is about and clicking here drops down a logical and well structured submenu with three main headings: inference for continuous data; categorical data; and measures of association. Beneath each heading is a manageable clutch of subheads, each of which leads to a side popup of specifics. The other entries (basic statistics; plots; power and sample size) are smaller affairs, there to support the main show.

Data is entered either into a worksheet or through an editor. Import of existing data from other programs covers a respectable range – apart from the usual suspects (ASCII, dBase, Excel, etc) it can be taken from other formats including BDMP, SAS, SPSS, Statistica, Systat, STATA. Export is more restricted, with only four options – ASCII, SAS, SPSS, STATA. Multiple datasets can be loaded and handled simultaneously, all accessed from the left-hand organiser pane. StatXact’s own data and results are stored in a set of retrievable file types.

Exactness

How exact is exact? Despite all the miracles which ubiquitous and ever more powerful computers have brought, they cannot yet do everything within a useful time. Ironically, their very speed has shortened our attention span – I would once have flattered and cajoled for weeks to get computer time on the old megalithic mainframe, before waiting days for the result of an analysis, but am now irritated if it takes more than a second or two. StatXact takes advantage of this power to do what it does, but there are limits. A true exact distribution is there whenever its calculation time is realistic – and, in turning over some pretty large data sets from the archive, I was impressed to see how often it is possible to get a true exact measure – but there are still times when you have to bow to the inevitable. For those occasions, an alternative strategy is pro-

What’s new

Though new to me, StatXact is in release six. For readers already familiar with the package from previous incarnations, here is a very quick summary of features new in this one:

- The efficient new project organisation-based interface, Cytel Studio, highlighted in the product name and mentioned in the main text;
- New exact procedures including confidence intervals, superiority, equivalence, non-inferiority, difference of proportion, odds ratio for paired binomial data; sample size calculations for binomial and multinomial data;
- The basic statistics module, which includes ANOVA, histograms, scatter plots, and bubble plots, multiple linear regression, t-tests;
- Stratified $R \times C$ exact tests for unordered, single ordered, double ordered data; and
- Improved import and export, including direct export to other statistical programs and import without specifying a file type.

vided in the form of Monte Carlo simulation within the distribution space of the particular problem under study. This is not exact, but in most studies it very definitely delivers a closer model than assumption of the theoretical distribution. Default settings for the Monte Carlo p-value run give a result around the 99 per cent accuracy level, but can be over-ridden to tweak the relative importance of accuracy versus execution time.

Using the program, you therefore have available three levels of ‘exactness’ – the true exact measure; a Monte Carlo approximation, which can be driven as close to exact as patience or time constraints allow; and a value derived from the theoretical distribution. In any case, since the exact distribution is asymptotic to the theoretical one as the sample grows, any retreat from exactness occurs only as its importance diminishes. Moving between the options is a painless business, and you remain in complete control of the balance.

For various unavoidable reasons, this was a short-term review conducted within the software’s evaluation period. I wasn’t, therefore, able to follow my usual routine of putting it to work for real on a live study. I did, however, have long enough to explore it thoroughly and also to try it out on some archived results (some my own, mostly those of colleagues so that my own defensiveness wouldn’t be an issue) from past studies using standard theoretical distribu-

tions. This last is not as easy as it sounds; the whole point of a package like this is that it does what is difficult to do by other methods, so checking up on it involves a mixture of assumption, cunning, and hours of tedious slog. Luckily, the benefits of the software are greatest in the case of small samples – and that is the area where manual checks are most feasible.

I isolated some cases where no analysis had been done because sample size was too small. To that I added one which only just scraped over the line – sample size small, just about meeting the various rule-of-thumb criteria, but importance high, so the analysis went ahead with suitable caveats. In the process of manually enumerating the larger set (though still small by all operational standards), I rediscovered the trauma of Dr P’s assignments. I found myself hoping that all this work would reveal some small fault with the software as a reward; but no, it really does provide an exact result to the level I have the patience to investigate.

Revisiting 180 old analyses

Altogether, I revisited a random sample of more than 180 old non-parametric analyses. In only one case did the exact value result throw serious doubt on the original finding (and that one was intuitively shelved as unsafe at the time). But that’s not the whole story, and is really beside the point. Far more interesting is the number of occasions where an exact value (or a Monte Carlo approximation to it) added useful additional insight. Knowing the divergence between theory and reality, there is always a tendency to build into statistical recommendations and decision-making a safety-margin based on gut feeling and risk assessment. In a significant proportion of instances, that safety-margin was shown to be much larger than necessary; in a few, it was only just enough.

Though they are the core issue (what I must these days learn to call the ‘unique selling point’), StatXact doesn’t restrict itself only to the exact non-parametric. There are distributionally asymptotic measures here, basic statistical descriptors, the plots menu, and so on. The package is a well-rounded candidate for the post of single statistics program in the right circumstances.

If non-parametrics are your concern, and everything else is subsidiary, this may well be your best choice. In a larger organisation using non-parametrics within a larger mix, it will add precision, muscle, and confidence to your existing set-up.