

GETTING REAL STATISTICS INTO ALL CURRICULUM SUBJECT AREAS – CAN TECHNOLOGY MAKE THIS A REALITY?

James Nicholson, Jim Ridgway and Sean McCusker
SMART Centre, Durham University, Leazes Road, Durham, DH1 1TA, UK
j.r.nicholson@durham.ac.uk

ABSTRACT

Technology has revolutionised society and it has revolutionised the way in which statistics, as a professional discipline, is done. The collection of data is growing exponentially both in relation to the quantity of data assembled on any particular measure and also in relation to the range of topics, and the measures, on which data is collected. Accessing data has become much simpler, and tools for exploring, manipulating and representing that data visually have multiplied, both in commercially available software and open-source freeware. However, the curriculum in schools in the UK is constrained by important factors which restrict the use of technology in assessment. The statistics curriculum is largely dull and does not address the core issues of most relevance in statistics today. Here, we explore ways in which technology can enhance the teaching of subjects in which statistics are used, and also the teaching of statistics within mathematics.

INTRODUCTION:

The Royal Statistical Society launched GETSTATS in 2010, a 10 year statistical literacy programme which addresses statistical literacy across the full spectrum of society - including formal education across many disciplines. They initiated a debate looking at the mathematical and statistical needs of all subject areas at school level and at undergraduate programme level, with an eye to preparation for involvement in research. This includes two separate strands looking at the needs and current state of provision within humanities and social sciences and separately within science, engineering and technology (STEM) subjects.

This initiative is partly in response to a number of major reports (ACME (2011 a, b), Hodgen, Pepper, Sturman and Ruddock (2010), Vorderman, Budd, Dunne, Hart and Porkess (2011)) which identify the UK as lagging well behind many other countries in terms of the number of 16 – 19 year old students who are studying some mathematics, with recommendations that have been accepted by the current government that all students should study some mathematics, pre-university. However, A-level mathematics (a specialist ‘gatekeeper’ course for STEM subjects for 16 - 18 year olds) is currently taken by nearly 80,000 candidates but is not an appropriate qualification for most of the rest of the cohort. Key questions arise: what should they study? and, crucially, who is going to teach them (since there is already a severe shortage of well-qualified mathematics teachers to deliver the current courses)?

Holmes (2000) analysed the statistical requirements embedded in all subjects which are available in the English National Curriculum; Porkess (2012) has undertaken an updating of that work. His report analyses the subject demands in terms of 4 components of statistical work:

- A. Problem analysis
- B. Data collection
- C. Data presentation
- D. Data analysis

Statistics is currently taught within mathematics. In primary school, (ages 5-11 years), components A, B and C are treated coherently at an appropriately simple level, but at GCSE (the high stakes assessment at the end of compulsory schooling, age 16) only component C has any prominence and at A-level (at age 18 / 19) only components C and D are involved. Psychology, Geography and Biology address all 4 components, yet anecdotal evidence about statistics within those subjects is often of it being treated as a black box of techniques, rather than as a way of thinking about the world.

One consequence of the curriculum specialisation possible in early secondary school in the UK is that pupils seem to split into two groups which almost seem to inhabit non-overlapping universes – the worlds of social science and of mathematics and statistics. The vast majority of students never encounter real social problems as substantial contexts in which they work to develop

or apply mathematical or statistical techniques; students who study mathematics and statistics do not encounter data where mathematics and statistics offer real insights into social issues. For students who work in social sciences and humanities, the academic work is substantially text based with occasional headline statistics to support an argument, but those headline statistics commonly summarise aggregated data, taking no account of the differences which exist between different sectors of the population.

The new curricula introduced in New Zealand (Ministry of Education, 2007) and South Africa (South Africa Curriculum, 2005) offer encouraging examples, with multivariate data being a core part of the statistical reasoning from an early age. There, young children do not find it strange that boys and girls behave differently; that children of different ages behave differently; that things change over time etc. In these curriculums, children look at data presented in an accessible form and in a context for which they have the language which allows them to be able to talk about the data sensibly; consequently they are able to make sense of the stories in the data.

We will identify some opportunities for ways in which the use of real statistics can enhance the teaching of subjects other than mathematics, and then look critically at the statistics curriculum within mathematics in the UK and suggest ways in which the use of technology might offer opportunities for improvement.

STATISTICS WITHIN SUBJECTS OTHER THAN MATHEMATICS

The use of technology offers opportunities to break that cycle of pupils splitting into non-overlapping universes – interfaces such as the ones we have been developing at the SMART Centre seem to sit very comfortably in a somewhat improbable niche – subject specialists in non-mathematical disciplines who feel anxious about data seem to lose sight of that anxiety when using the interface because what they and the students see are the subject content stories in the data; and mathematics teachers who do not have a lot of discussion in their classrooms seem to feel more confident that they can moderate discussion on this sort of data – perhaps perversely because the stories in the data are more complex than the data they encounter in traditional data exercises and therefore it is clear there is no one right answer. In both cases this seems to be liberating for the students, but also seems to make teachers more comfortable about the possibility of discussion.

We have some datasets from physical and life sciences, such as oxygen production and epidemiology, where the use of multivariate data can provide interesting insights into biological, chemical and physical processes. In social science arenas the use of real data sets on topics central to the curriculum offers opportunities to address important social science ideas, such as adopting a critical stance to the provenance of data – which are much more real to students when presented in concrete case studies rather than in an abstract list of principles. In education data, measurement and categorisations include: did the pupil pass 5 GCSEs at grade or above, including English and Mathematics?, gender (easy to categorise as boys and girls isn't it?), eligibility for free school meals, ethnicity ... (see Nicholson, Ridgway & McCusker, 2011).

We think the niche our tool sits in is important and we are not aware of other tools which offer anything quite like it: it works with multivariate data, but it is accessible to naïve users in that it is based on comparative bar charts and so the structure of the information is familiar and recognisable to almost all users. The power is in the user's ability to manipulate the display, using sliders to look at different sections of the population, swapping the position of variables in order to make explicit different comparisons, and in the capacity to display related data sets in separate tabs within the one file. This gives the user the capacity to actively explore the data and identify the stories in the data for themselves.

The SMART centre interface is much less sophisticated than *Gapminder* or *eXplorer* in terms of the variety of types of data it displays, but it can be used to facilitate data exploration by a much wider audience. It addresses a very different area of statistics than data analysis software does – *Fathom* and *Tinkerplots* are powerful educational environments which allow users to construct multiple representations of multivariate data, but these representations are dealing with individual data cases where the SMART centre interface is displaying aggregated data (similar to some of the data sets displayed in *Gapminder* or *eXplorer*). *Fathom* and *Tinkerplots* can be used by skillful teachers to help students build a sound understanding of many statistical concepts, but they also require a considerable investment of time in learning to drive them effectively. Other standard

platforms such as *Autograph*, *Minitab*, *R*, *SPSS*, *FxStat* or the RSS Centre for Statistical Education's *Data Tool* provide options which are appropriate for particular user groups but they do not allow the exploration of the relationships within aggregated data that the SMART centre's interface is designed for. Developing the interface so that a user can analyse a large data set to produce their own *n*-way table of aggregated data and be able to import that directly into the SMART centre's interface would significantly enhance its usefulness. Currently users are limited to the pre-constructed representations, and choice of variables, built in by the author team for that particular data set.

The use in social science subjects and in physical science subjects have somewhat different characteristics. An example taken from each area follows:

Statistics within social sciences

In early August 2011 there was a 4 day period with serious public disorder in a number of English cities. There was considerable public debate about the reasons for the disorder and also about the profile of those involved in the disorder. The Ministry of Justice (2011) published a special Statistical Bulletin on September 15th 2011, which generated further public debate with considerable political commentary from all sides. This sort of data release offers a huge opportunity for a focus on real social issues, and on the role of mathematics and statistics in understanding the nature of the issues, and can be made available for use in schools in an accessible medium.

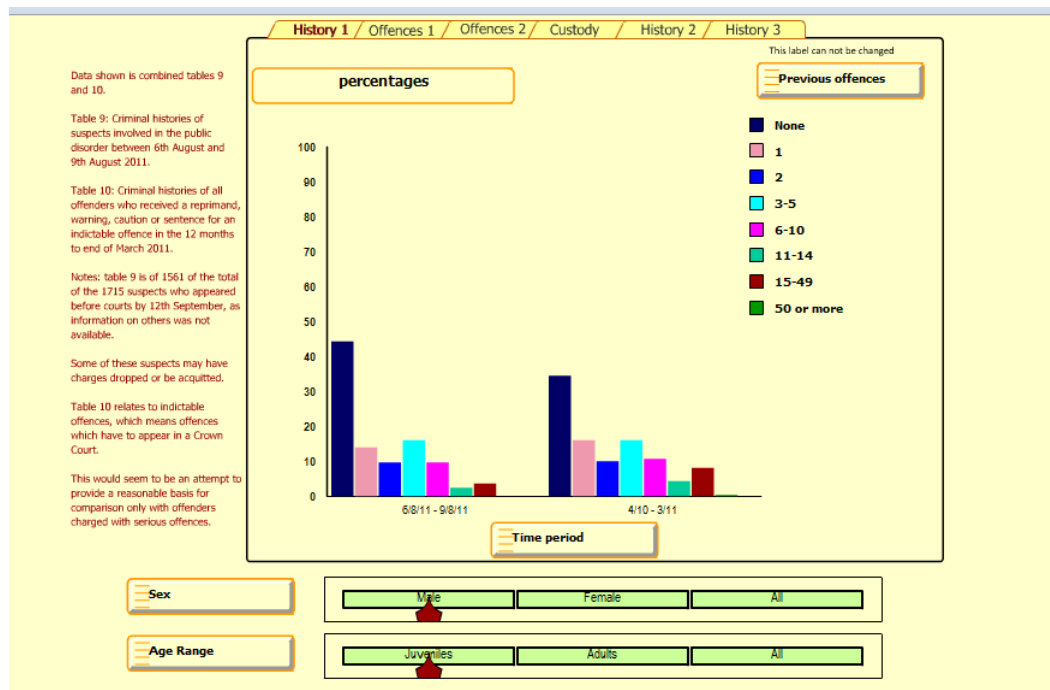


Figure 1: screen dump of public disorder data visualisation

Figure 1 shows a comparison of the criminal histories of male juveniles during the public disorder and during the 12 months up to March 2011. The sliders allow users to look at females or both genders combined and also adults only or all ages together. There are a number of related data sets accessible by clicking on the tabs at the top. This data set is relevant to a range of social science disciplines: media studies, politics, sociology, citizenship, history and psychology would all have immediate and substantive interest in it from their different perspectives.

Because the interfaces are intuitive to use, teachers seem comfortable in engaging with a new data set and we see a real opportunity to change the perception that a lot of students and teachers have that statistics is boring and irrelevant.

The data display is generic, and is not tied to a particular age or stage of statistical sophistication. It is usable by adults who are statistically naïve, but who are interested in the subject matter. It is easy to distribute such displays electronically, to add them to media web pages, or to base *YouTube* videos on the content. These features make it possible to blur artificial barriers between formal education and everyday life. The facility does, however, place an onus on the provider to exert extreme care – specifying sources, explaining metadata, and adding appropriate cautions about interpretation, and pointing to key statistical ideas (here, for instance, one might ask about systematic bias – why do so many of the people arrested already have a police record? It seems likely that it is far easier to identify and arrest miscreants filmed on CCTV if their photographs (and addresses) already exist on police files than if they do not). While these notions might be the explicit focus of attention in one subject – for example in media studies looking at the ways media reporting or the use of social networking media may actually help to shape the events, often they will not be, but understanding the distorting effect of a systematic bias will be important in trying to interpret the information displayed by the data.

Ridgway, Nicholson & McCusker (in press) discuss the challenges and opportunities of the ubiquitous access to sophisticated data displays more fully, but it is important to note that there is substantial cause for concern about the variability in quality and how (or even whether) users can identify meaningful and reliable data amidst an almost overwhelming volume of possibilities that internet search engines will provide.

Statistics within physical sciences

Photosynthesis is a chemical process by which carbon dioxide and water react in the presence of light to produce biomass, with oxygen as a by-product. The level of oxygen production (a surrogate measure for the rate of photosynthesis) varies with the amount of sunlight, temperature and with the concentration of carbon dioxide in the air, giving three factors whose effects can be explored. These interact with one another, as well as displaying behaviour which cannot be expressed in terms of the functions that students are familiar with. Restricting their experience to a world in which only behaviour which behaves according to polynomial, exponential or other simple functions does them a grave disservice. For example, there is a minimum light level required before any oxygen is produced, and there is a limiting light level beyond which the production rate plateaus. For any given light level there is an optimum temperature at which the rate of production of oxygen is a maximum; and as the concentration of carbon dioxide in the air increases, the light level at which the plateau effect is apparent increases. In Figure 2 below the screen dump on the left shows a lower concentration of carbon dioxide in the air than the one on the right (the carbon dioxide concentration level is on the lower slider). The interaction between the concentration of carbon dioxide and light intensity is illustrated very nicely. Students can articulate the relationships very well, even though they do not have a grasp of the mathematical functions that describe oxygen production.



Figure 2: screen dumps of oxygen production in photosynthesis visualisation

There are important differences between understanding a physical (largely deterministic) function which can be explored by repeated experiments, and understanding social science data which often comprises very large data sets derived from observations. In the case of social science data students may need to identify the behavioral characteristics of different groups as a precursor to developing theories which explain those patterns of behaviour – which may be sociological or psychological theories or located within some other branch of social science.

Students moving on to college will encounter multivariate data in many subjects, and be expected to carry out statistical analyses using software. However, currently few if any of them have any mental image of what a multivariate data ‘looks like’ or what an interaction between variables means in practical terms. Worse, most statistical packages assume linearity; many physical phenomena (e.g. photosynthesis) show non-linear relationships (there is a real risk that dependence on SPSS, without an awareness of the assumptions about linearity and normality made in many of the analyses that are available, has hidden the non-linear relationships in social science phenomena from both theorists and policy makers). We believe that experiencing data contexts that students are familiar with in these sorts of interfaces may help them develop a schema in which to reference developing statistical concepts.

We have evidence (see Nicholson, Ridgway & McCusker, 2009) that children as young as 12 years who are at the lower end of the academic spectrum can identify the main patterns in data (including cases where there are anomalies in otherwise consistent patterns), when they have access to data in interactive displays. Even younger children can work meaningfully with multivariate data displays that they have created in the form of tactile graphs built from LEGO® bricks (du Feu, 2005).

STATISTICS WITHIN MATHEMATICS

Porkess (2012), in a report commissioned by the Royal Statistical Society and the Institute and Faculty of Actuaries on the future of statistics in UK schools and colleges, recommends that statistics remains in the mathematics curriculum. However, there is little in his report which has not already been advocated, and the report represents a major missed opportunity to improve statistics education. The current curriculum content is largely dull and does not address the core issues of most relevance in statistics today. There is a major emphasis on mechanical calculation of summary statistics and production of graphical displays that are always automated in the real world, and almost no emphasis on the critical evaluation of what are appropriate measures or displays, nor of how to decide what should be measured.

Ridgway, Nicholson & McCusker (2006) reported on the small proportions of the assessment credit in A level statistics modules for any form of modeling or interpretation. McCusker, Ridgway & Nicholson (2010) and Nicholson, Ridgway & McCusker (2009) explored ways in which poor assessment of statistics at GCSE level damages the image of statistics via the use of inappropriate contexts. Regrettably, although new specifications have been put in place since those papers were written, the assessment does not appear to have improved.

The Porkess report makes no mention of the potential use of technology to make the study of statistics more accessible, as well as more in line with practice in the world of work whenever summary statistics or graphical representations are required. Risk has been brought in, but is treated almost exclusively as a context for probability calculation rather than risk being conceptualised as a combination of likelihood and the magnitude of the consequences. Crucially, ‘disasters’ are very often the end result of a sequence of events and much of risk management is about what steps should be taken as and when things start to depart from what is expected. The simplest example of this is in statistical process control (SPC), yet the normal methods for teaching SPC do not make use of technology to simulate a developing portfolio of information available to the operator on which they have to make a series of decisions as to whether or not to make an intervention, and if so, what the intervention should be.

The ways in which technology can enhance the teaching and learning of statistics are many and varied: it can enable students to work with large data sets; it can enable more time to be spent on critical evaluation of what the most appropriate representation of data is in either summary statistics or graphical form. The use of data visualisation tools would enable students to begin to develop mental models of the properties that multi-variate data sets can possess, and what the

important features are i.e. where are the large effect sizes, are there interactions etc.; animations of procedures can allow students to work through material in their own time to consolidate their understanding; simulations can help develop sound conceptual understanding of difficult topics – when they are used with good pedagogy. Much of this has been available for some time – it is much easier to produce animations now, and accessible visualisations of multi-variate data are relatively recent.

The website <http://understandinguncertainty.org/> makes available some tools which can be used to help students get to grips with some aspects of risk, but these fall short of providing good curriculum materials which could form the basis of a coherent contribution to a standard mathematics course. The 1996 IASE roundtable in Granada focused on *The role of technology in teaching and learning statistics*: the proceedings are available to download at http://www.dartmouth.edu/~chance/teaching_aids/IASE/IASE.book.pdf and it is instructive to look back at the collection of papers and see where the subject area was then, and how far things have progressed. There is also an extensive literature on various aspects of the use of technology which can be found in the proceedings of ICOTS, IASE and ISI conferences as well as in statistics education journals (see for example Hancock, Kaput & Goldsmith (1992), deMas, Garfield & Chance (1999), Wild & Pfannkuch (1999), Erickson (2001, 2002), Chance (2002), Nicholson, Mulhern & Hunt (2002), Shaughnessy & Pfannkuch (2002), Finzer & Erickson (2004), Chance & Rossman (2006), Grunewald & Mittag. (2006), Chance, Ben-Zvi, Garfield, and Medina (2007), Finzer, Erickson, Swenson, and Litwin (2007), Harradine & Baker (2008), Konold (2010)). It is interesting to note that the underlying pedagogical principles are not often the primary focus of attention of these articles, and the understanding by classroom teachers of the important pedagogical principles is generally weak.

Uses of technology to improve teaching

Technology can be used innovatively by teachers to illuminate concepts which are beyond the mathematical competence of pupils to prove at their current stage of mathematical development. Many simulations have been around for some time, and are freely available, but the pedagogy of their use is not widely articulated and many teachers are much less able to critically evaluate electronic resources than text-based resources, and do not easily identify sound pedagogical strategies for their use in a classroom. This failure to articulate appropriate pedagogy on the part of statistics educators who produce these resources may be at least a partial explanation for why they are not more widespread in their use.

We offer three examples to illustrate this. These are at the opposite end of the spectrum to the use of the SMART centre's interface described in the first section which was accessible to naïve users with no sophisticated knowledge of statistics. The content would be covered in A-level mathematics or statistics courses or at university in the UK. These simulations run in Excel, and were embedded in a web environment which allowed users to explore related concepts. Each simulation appeared in the top half of an Excel worksheet, and the lower half contained some guided exploration activities to help users work with the simulation in the most effective way to build sound conceptual understanding.

1. Regression and correlation

The concepts of correlation and of the line of regression are often misunderstood, and calculations are judged to produce population parameters rather than being the statistics of a particular sample data set. Nicholson & Mulhern (2000) explored the confusion between deterministic and stochastic behaviour that students experience: so for example, the expectation of a random variable is not clearly recognised as deterministic because probabilities are used to calculate it and a regression line is often thought of as THE relationship (i.e. thought of as deterministic) between the two variables rather than an estimate of it based on a particular sample of data – with the precision of that estimate being affected by the number of observations and on the strength of the underlying relationship between the two variables.

Figure 3 shows a screen dump from the DISCUSS regression materials (Hunt, Tyrrell & Nicholson, 2000a) which illustrates how the line of regression is dependent on the sample of data chosen from a bivariate population. The example displayed shows a sample of 15 points

highlighted in red chosen randomly from a larger population (non-sampled points are shown in grey), with the line of regression for that sample of points shown in red. The lines of regression from previous samples are shown in yellow and the blue line is the line of regression for the whole data set. Running the simulation, a teacher can explore the extent of the variability in the slope and position of the regression lines for different samples, and by changing the sample size, they can explore with a class how the stability of the parameters of a line of regression depends on the size of the sample used to construct the regression line, and how distance from the centre of the data set affects the likely accuracy of the prediction. Extending this to a number of underlying data sets would allow users to also explore how the underlying correlation between variables in the population affects the stability of the prediction.

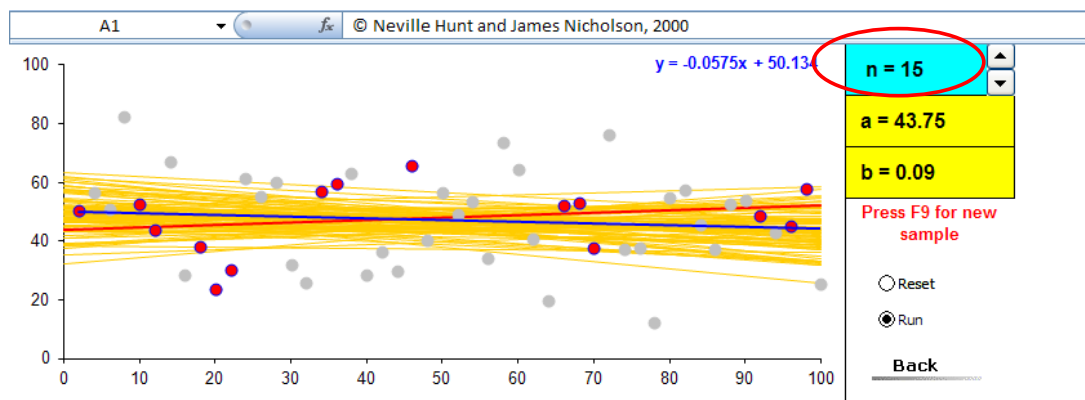


Figure 3: screen dump to illustrate the stochastic nature of the line of regression

2. The Central Limit Theorem (CLT)

Technology can offer an approach that is a little more subtle than the treatment often given in textbooks. The CLT says that the distribution of sample means becomes approximately Normal as the sample size increases, no matter what the underlying population is, and at school level an arbitrary cutoff of $n = 30$ is applied indiscriminately. However, for a uniform distribution, or for a symmetric binomial distribution [and for almost any reasonably symmetric population] the distribution of sample means is approximately Normal with rather smaller sample sizes than 30.

There are two important facts which get lost sight of very often:

- The CLT says nothing about the standard error of the mean – that result is derived from algebraic manipulations of the definitions of mean and variance – it is solely about the probability distribution of the sample mean.
- For any underlying population, increasing the sample size will make the sampling distribution of the mean look more like the Normal (except when the underlying population is itself Normal, in which case the sampling distribution of the mean will also be an exact Normal).

Figure 4 shows screen dumps from the DISCUSS estimation materials (Hunt, Tyrrell & Nicholson, 2000b) which illustrate how the behaviour of the sampling distribution for different distributions and for varying sampling sizes can be explored qualitatively. The top screen dump is the sampling distribution for samples of only 5 from a uniform distribution, but the shape of the Normal distribution is already beginning to emerge. In contrast, the samples of size 5 from an exponential distribution shown in the middle screen dump still display the skewed property of the underlying population. When the sample size is increased to 20, as shown in the bottom screen dump, the skewness has almost disappeared and the Normal distribution is beginning to emerge, though the interval groupings used in the display do not allow the detail to be seen as clearly as one would like.

There are a lot of places where statistics uses rather arbitrary cutoffs. The reality is that there is a continuum on which you have to make a decision at some point that you will adopt a different stance – for example to move from accepting a null hypothesis to rejecting it, because a probability has moved from just above 5% to just below it.

The simulation lets you choose different populations and different sample sizes, which is no more and no less than what you need to explore the CLT. The accompanying activities encourage the user to explore methodically the way the sampling distribution changes as the sample size increases – for one underlying distribution first, and to make a judgment as to where an approximation by the Normal would be reasonable. The underlying distribution is then changed, and the process is repeated, so that students can explore the extent to which the underlying distribution makes a difference to the sample size needed for the Normal to be an acceptable approximation.

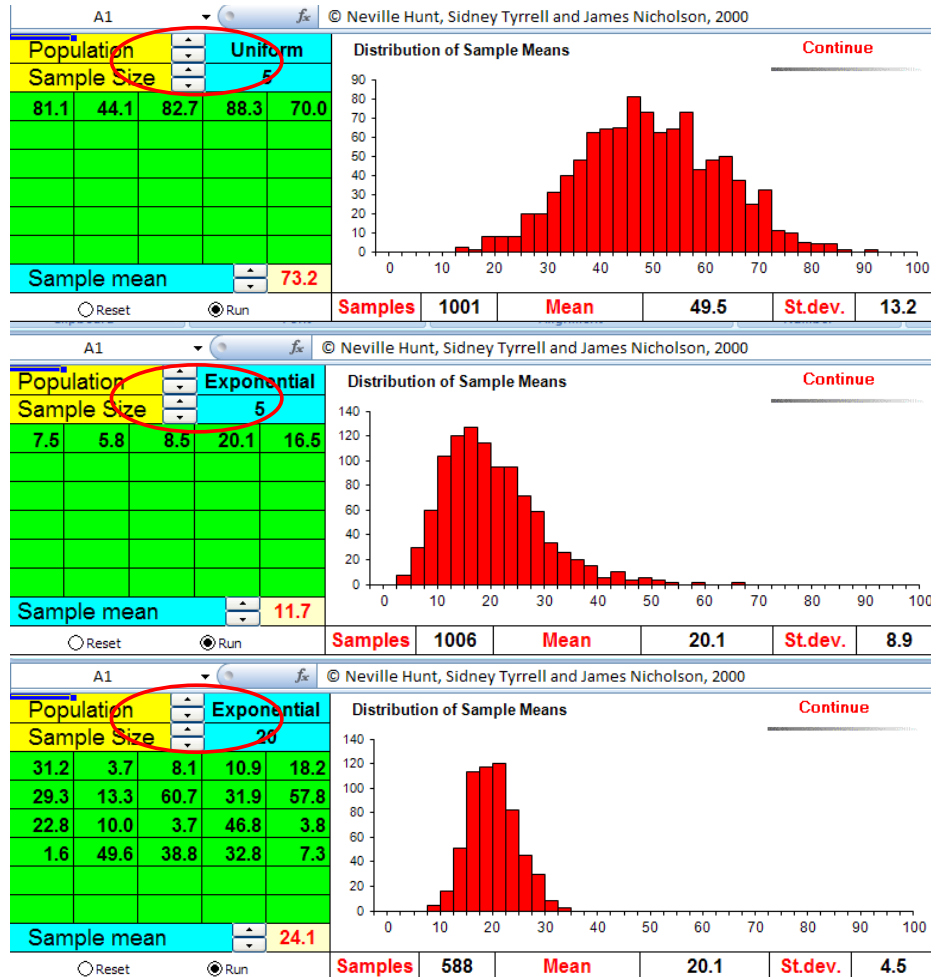


Figure 4: screen dumps of sampling distributions to illustrate the central limit theorem

3. The *t*-distribution

Textbooks deal with the mechanics of using the *t* distribution admirably in most cases and advanced texts show the functional form of the distribution, but there is little commentary as to why the family of *t* distributions compensates for that extra uncertainty introduced by having to estimate the variance from a small sample. In seeking to understand the effect of estimating the variance, the key questions are: how does the estimator of the variance behave? And how does that behaviour affect any confidence interval we construct based on that estimator?

Figure 5 shows two screen dumps from the DISCUSS estimation materials (Hunt *et al.*, 2000b). In the top screen dump, the sample size is 5, the standard deviation of the population being sampled is actually 15, but the radio button chosen on the bottom left indicates that the sample data is being treated on the basis that the standard deviation is not known (and is to be estimated from the sample). The grey column of figures shows the estimated standard deviations for 10 different

samples, which range from 3.7 to 24.4. The lower screen dump shows 10 samples of size 15 and now the standard deviation estimates are much more consistent, ranging from 13.4 to 20.3. The effect of the different estimated standard deviations on the confidence interval width within each screen dump is dramatically evident. By running the simulations a number of times one can observe that the standard deviation is underestimated more often than it is overestimated, but by less on average than when it is overestimated. Consideration of what effect this will have on the likelihood of capturing the true population mean if the Normal distribution cutoff for a 95% confidence interval is used will lead to the conclusion that it will be much less than 95%. This is only a qualitative discussion of something where the quantitative reasoning depends on mathematics which is more sophisticated than is available at this level, and it provides a powerful rationale not only for why the cutoffs in the t distribution are different to the cutoffs in the Normal distribution. The evident difference in variability of estimates as the sample size varies offers a rationale for the existence of the family of t distributions, and why the cutoffs decrease towards the corresponding Normal cutoff as the sample size increases.

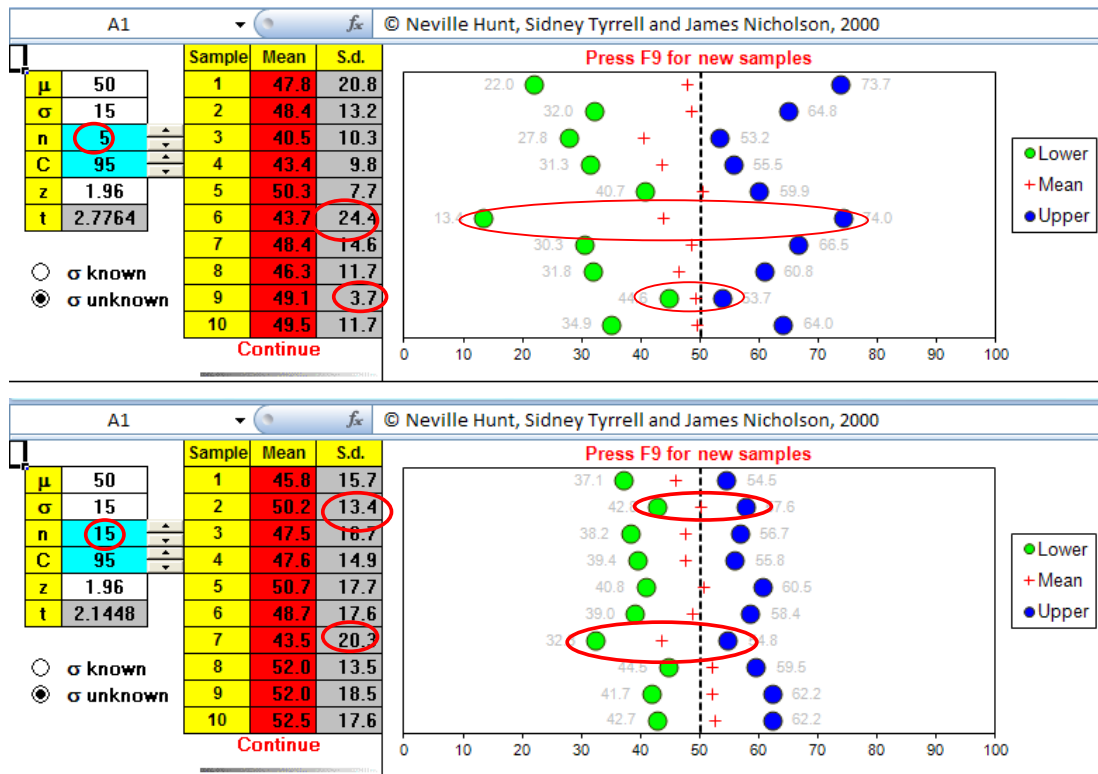


Figure 5: screen dumps of t distribution simulation of confidence intervals

In all of these three examples, one key aspect is that the simulations allow a difficult concept to be illuminated visually, with opportunities for the subtleties to be explored in a relatively short time frame, but it does require teachers to have a firm grasp of the concepts and to know how to use the simulation to draw out the important features.

DISCUSSION

The use of technology enables large data sets to be managed. We have discussed two very different types of large data sets which we believe have the potential to enhance the teaching and learning of statistics: firstly the use of data visualisations with large data sets in social sciences and physical sciences to provide an accessible data landscape, and secondly the use of large data sets generated in simulations to understand difficult statistical concepts.

If we are serious about increasing the exposure of our students in the UK to statistics within the curriculum, there needs to be a recognition that there is not currently, nor is there likely

to be in the foreseeable future, the capacity for this to be undertaken by mathematics teachers. There is a substantial shortfall in the number of specialist mathematics teachers required in most areas of the UK. If we argue that statistics needs to be embedded within other curriculum areas, then we need to work with other subject areas to develop curriculum materials which introduce worthwhile statistical ideas in a way which is manageable for the average teacher in those other subject areas. These data visualisations allow the user to concentrate almost exclusively on processing the information available in the subject context, because the manipulation of the interface is so transparent. While many users do not see it as doing any mathematics or statistics, because they do no calculations or graph constructions for themselves, they are intuitively distinguishing between large and small effect sizes, identifying interactions, exploring the effects of disaggregation, and generating discussions about the reliability of the data – in other words they are communicating a large number of the big ideas in statistics, but without the focus on the minutiae of statistical techniques which seem to be threatening to many students, and teachers in other disciplines.

There is a very strong argument for moving away from hand calculations and constructions towards automation of calculations of summary statistics and graph production, and concentrating more effort on the interpretation of graphs and on identifying what the appropriate summary statistic or graph should be in a given context. Increasing the use of technology in this area would open opportunities for the same technology to be used to support conceptual development through the use of animations and simulations as teachers become more confident in its use, but there is work to be done to provide appropriate pedagogical support for teachers seeking to introduce technologies into their classroom.

The potential rewards to the statistics education community of creating interactive data displays based on topics of urgent social concern are high. The engagement of social scientists with statistics, and the engagement of mathematicians in debates about social issues – and public debate grounded in evidence, should generate higher levels of statistical literacy in the population as a whole, and may help to reduce the fear many people currently associate with numbers (see Ridgway *et al.*, 2011 for a fuller discussion on developing statistical literacy).

Teachers seem to be well aware that there is a difference in the type of thinking required to successfully cope with A level Statistics within Mathematics. They are much less able to articulate what that difference is, and particularly what the source of difficulties encountered by their students is. Since the procedural / computational side of statistics is becoming increasingly automated, it is essential that the pedagogical issues relating to statistical thinking at this level are explored properly, and that much more detailed work on the pedagogy of using technology to support the development of statistical thinking is required, and made accessible to classroom teachers.

ACKNOWLEDGEMENT

Some of this research was funded by the Nuffield Foundation Grant EDU/33713. The views expressed here are entirely those of the authors.

REFERENCES

- ACME (2011a) *Mathematical needs: Mathematics in the workplace and in Higher Education*. http://www.acme-uk.org/media/7624/acme_theme_a_final%20%282%29.pdf
- ACME (2011b) *Mathematical needs: The Mathematical needs of learners* http://www.acme-uk.org/media/7627/acme_theme_b_final.pdf
- Chance, B. & Rossman, A. (2006) *Using simulation to teach and learn statistics*. In A. Rossman and B. Chance (eds.). Working Co-operatively in Statistics Education. Proceedings of the Seventh International Conference on Teaching Statistics. International Statistics Institute, ISBN-10 90-73592-25-9.
- Chance, B. (2002). Components of Statistical Thinking and Implications for Instruction and Assessment, *Journal of Statistics Education*, vol 10(3).
- Chance, Beth; Ben-Zvi, Dani; Garfield, Joan; & Medina, Elsa. (2007). The Role of Technology in Improving Student Learning of Statistics. *Technology Innovations in Statistics Education*, 1(1). Retrieved from: <http://escholarship.org/uc/item/8sd2t4rr>

- delMas, R., Garfield, J. & Chance, B. (1999) A Model of Classroom Research in Action: Developing Simulation Activities to Improve Students' Statistical Reasoning *Journal of Statistics Education* 7(3)
- du Feu, C. (2005) Bluebells and Bias, Stitchwort and Statistics. *Teaching Statistics*, 27(2), 34-36
- Erickson, T. (2001). *Data in depth: Exploring mathematics with Fathom*. Emeryville, CA: Key Curriculum Press.
- Erickson, T. (2002). *Fifty Fathoms: Statistics demonstrations for deeper understanding*. Oakland, CA: EEPS Media.
- Finzer, W. & Erickson, T. (2004) *Curriculum Innovations Based on Census Microdata: A Meeting of Statistics, Mathematics, and Social Science*. IASE Roundtable: Curricular development in Statistics Education. Retrieved from: http://www.stat.auckland.ac.nz/~iase/publications/rt04/4.4_Finzer&Erickson.pdf
- Finzer, W., Erickson, T., Swenson, K. & Litwin, M. (2007). On Getting More and Better Data Into the Classroom. *Technology Innovations in Statistics Education*, 1(1). Retrieved from: <http://escholarship.org/uc/item/09w769>
- Gapminder <http://www.gapminder.org/downloads/>
- Getstats <http://www.getstats.org.uk/>
- Grunewald, W., and Mittag, H-J. (2006). The use of advanced visualisation tools for communicating European data on earnings to the citizen. In A. Rossman and B. Chance (eds.). Working Co-operatively in Statistics Education. Proceedings of the Seventh International Conference on Teaching Statistics. International Statistics Institute, ISBN-10 90-73592-25-9.
- Hancock, C., Kaput, J. J., & Goldsmith, L. T. (1992). Authentic inquiry with data: critical barriers to classroom implementation. *Educational Psychologist*, 27(3), 337-364.
- Harradine, A. & Baker N. (2008, July). Birthing big ideas in the minds of babes. In C. Batanero, C. Reading & A. Rossman (Eds), *Joint ICMI / IASE Study: Teaching Statistics in School Mathematics. Challenges for Teaching and Teacher Education. Proceedings of the ICMI Study 18 and 2008 IASE Round Table Conference*.
- Hodgen, J., Pepper, D., Sturman, L., Ruddock, G.(Nuffield Foundation) (2010). *Is the UK an outlier? An international comparison of upper secondary mathematics education*. www.nuffieldfoundation.org/uk-outlier-upper-secondarymaths-education
- Holmes, P. (2000). *Statistics Across the English National Curriculum*. Royal Statistical Society Centre for Statistical Education, <http://www.rsscse.org.uk/resources/natcur.html>
- Hunt, D. N., Tyrrell, S. & Nicholson, J.(2000a). *DISCUSS Regression and Correlation*. <http://nestor.coventry.ac.uk/~nhunt/regress/>
- Hunt, D. N., Tyrrell, S. & Nicholson, J.(2000b). *DISCUSS Estimation*. <http://nestor.coventry.ac.uk/~nhunt/sample/>
- Konold, C. (2010). *The Virtues of Building on Sand*. Plenary lecture at ICOTS8 – available as video lecture from http://videlectures.net/icots2010_konold_tvbs/
- McCusker, S., Nicholson, J. & Ridgway, J., (2010). Statistics Assessment: The Good, The Bad and The Ugly. In C. Reading (Ed.), *Data and context in statistics education: Towards an evidence-based society. Proceedings of the Eighth International Conference on Teaching Statistics (ICOTS8, July, 2010), Ljubljana, Slovenia*. Voorburg, The Netherlands: International Statistical Institute. www.stat.auckland.ac.nz/~iase/publications.php
- Ministry of Education (2007). *The New Zealand curriculum*. Wellington, New Zealand: Learning Media Limited.
- Ministry of Justice (2011). Statistical bulletin on the public disorder of 6th to 9th August 2011. Downloaded from: <http://www.justice.gov.uk/downloads/statistics/criminal-justice-stats/august-public-disorder-stats-bulletin.pdf>
- Nicholson, J. R. & Mulhern, G. (2000) *Differences in Mathematical and Statistical Thinking: Interpretation and Uncertainty – a discussion paper*. Psychology of Mathematics Education Conference 24, Hiroshima, Japan. Available at: <http://www.dur.ac.uk/smart.centre/publications/>
- Nicholson, J. R., Mulhern, G. & Hunt, D. N. (2002). Wizardry or Pedagogy? What is the driving force in the use of the new technology in teaching statistics? *Proceedings of ICOTS VI*. Voorburg, the Netherlands: International Statistical Institute.

- Nicholson, J., Ridgway, J., & McCusker, S. (2009). *One small step for a pupil – one giant leap for citizens*. In P. Murphy (Ed.), Proceedings of the International Association for Statistics Education Satellite Conference, Next Steps in Statistics Education. Durban SA: International Association for Statistical Education. Online: www.stat.auckland.ac.nz/~iase/publications.
- Nicholson, J., Ridgway, J., and McCusker, S. (2011). Visualise then Conceptualise. *Social Science Teacher*, 40(3), 8-12.
- Open Statistics eXplorer <http://ncva.itn.liu.se/explorer/openexp?!=en>
- Porkess, R. *The Future of Statistics in our School and Colleges*. The Royal Statistical Society and the Actuarial Profession, London. ISBN 978-0-9571305-0-0.
- Ridgway, J., Nicholson, J, and McCusker, S. (2006) Teaching statistics - despite its applications. *Teaching Statistics*, 28(2), 44-48.
- Ridgway, J., Nicholson, J., and McCusker, S. (2011). Developing Statistical Literacy in Students and Teachers. In C. Batanero, G.Burrill, and C. Reading (Eds.). *Teaching Statistics in School Mathematics-Challenges for Teaching and Teacher Education*. Springer: Dordrecht. New ICMI Study Series, Vol. 15, 311-322.
- Ridgway, J., Nicholson, J., and McCusker, S. (in press). *The Semantic Web demands 'New' Statistics*.
- Shaughnessy, J.M. & Pfannkuch, M. (2002) How faithful is Old Faithful? Statistical thinking: A story of variation and prediction. *The Mathematics Teacher* 95, 252 – 259.
- SMART Centre <http://www.dur.ac.uk/smart.centre/>
- South Africa Curriculum (2005). Online: www.education.gov.za/Curriculum/Curriculum.asp.
- Vorderman, C., Budd, C., Dunne, R., Hart, M., and Porkess, R. (2011) *A world-class mathematics education for all our young people*. Downloaded from <http://www.tsm-resources.com/pdf/VordermanMathsReport.pdf>
- Wild, C. & Pfannkuch, M. (1999). Statistical Thinking in Empirical Enquiry. *International Statistical Review* 67 (3) 223–248