Quarterly Newsletter

Volume 15, Issue 2, November 2015

Tails of Woe

by Felipe Dimer de Oliveira

It is often believed that arcane topics in advanced mathematics have their place only in the rarified echelons of academia. We forget that many of the technological advances that we now take for granted were once the preoccupation of bright minds with too much time on their hands. Insurance is no different: for over 400 years it has brought many mathematical innovations to practical use. In this article we explore how an understanding of fat-tailed distributions has applications to insurance and risk management.

Statistics is the art of making sense of the world by analysing and interpreting seemingly random data. Large amounts of information such as the height of every Australian or the income of families in NSW can be described by a distribution expressed as a variation around an average value. It is sometimes overlooked, however, that such a representation of reality is not reality itself: it can help us to understand the world, but not tame it.

The search for order in data via statistical analysis is usually underpinned by the idea that every process is somehow constrained by nature. In the case of heights of people, for example, there are physical limits to the height a human can reach which are dictated by anatomy and physiology. These boundaries show themselves as a bell-shaped curve around a well defined average value. In layman's terms, this means that, if we make an estimate of the height of some unknown Australian as the average value for the population, we cannot be very far off the truth. There are no adult humans 10 meters in height, for example.

There are, however, exceptions to the normally distributed attributes described above; some processes seem unconstrained by any physical limits or typical scales. A classic example is the distribution of populations of human settlements; these can vary from a few hundred to tens of millions - a variation of several orders of magnitude. Instead of a bell-shaped curve, its distribution shows what is called power-law behaviour.

Without getting into the technicalities, the main property of data following a power-law distribution is its lack of scale. For the size of cities, for example (more specifically US cities (Newman 2006)), there are a hundred times more cities with populations of 100 thousand than there are with 1 million inhabitants, and a hundred times more cities with 10 thousand than those with 100 thousand inhabitants. This pattern of a 100-fold decrease in the number of cities for every 10-fold increase in population repeats itself over many orders of magnitude. Such data when plotted in a log-log graph will appear as a straight line with angular coefficient 2 (as in $100=10^2$): this is the main coefficient that characterises a power law.

Even though we can calculate an "average city size" from the data shown in Figure 1, such a value is hardly representative of most cities. This is also true of some catastrophe losses (Kousky and Cooke 2009). In fact in the absence of more detailed modeling, reinsurers often assumed that catastrophe losses might follow a Pareto curve, which is one example of a power law.

A particularly interesting property of power laws is their "fat tails". Fat tails are often defined in terms of their skewness; in other words, they are asymmetric. More relevant for CAT losses is the fact that some fat-tailed distributions, in particular power laws with coefficients less than 2, lack a stable average. This means that, in the case of CAT losses again, the accumulation of yearly losses for the calculation of a long-

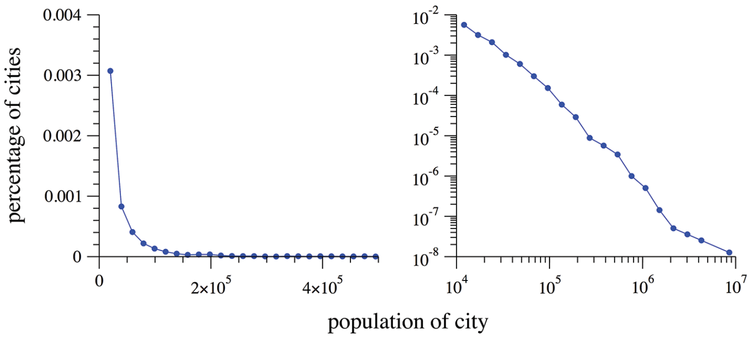

Figure 1: Histogram of the populations of all US cities with more than 10,000 inhabitants (left). On the right: the same data plotted in a log-log graph (source: Newman (2006)).

This Issue

- Tails of Woe
- QuakeNZ
- Auctions and the Government's Emissions Reduction Fund

Sponsors

Aon Benfield

Australian Reinsurance Pool Corporation

Guy Carpenter

IAG Insurance

QBE

Suncorp Group

Swiss Re

MACQUARIE University SYDNEY·AUSTRALIA

term average annual loss (AAL) will never converge to a stable value, and that a single event can shift the historical average significantly. This is in contrast with non-CAT losses, such as the annual tally of Australian car accident fatalities, which never departs far from the annual average of approximately 1,300 in any year.

Evidently, finite datasets will have a finite average: however, it can be demonstrated that, even in this case, averages from data drawn from fat-tailed distributions converge much more slowly than from non fat-tailed distributions. We will come back to this point in later discussion.

The implications of such behaviour to the insurance industry have been little studied - in particular, the mechanism by which fat-tailed distributions emerge and their consequences for accumulation.

The mechanism of a stochastic process is, broadly speaking, a mathematical description of the small-scale behaviour of a process. CAT-modellers are familiar with dealing with stochastic processes whenever they provide details on how to calculate losses for every individual address in a portfolio. This is the closest any CAT-modeller can come to translating the real world into the language of mathematics. Many will be familiar with the experience of changing the vulnerability or other small-scale parameter in a loss model only to find that the simulated loss for a particular event or footprint is completely unreasonable. This is the case because the processes that result in fat-tailed distributions are usually multiplicative (Mitzenmacher 2003), and even small changes in the underlying process tend to combine to produce large fluctuations in the final result.

This is probably one of the most challenging aspects of CAT-modelling. Ideally, it should be possible to infer the behaviour of the small parts by observing the whole: some phenomena, however, including natural catastrophes, show emergent properties in which the relationship between the whole and its parts is not at all obvious. Many generative models for power laws and some fat tails have been proposed and include the use of graph theory to monkeys typing randomly (that is not a joke!): very few attempts have been made to apply such models to NAT CAT problems.

The processes underlying non-CAT losses, on the other hand, can be easily understood with the help of the central limit theorem (CLT). The CLT states that the summation of many identically distributed random variables will always result in a well-behaved, bell-shaped, normal distribution. Thus, the distribution of annual losses resulting from car accidents mentioned above can be thought of as the summation of random variables representing the loss of each individual event. It is clear here that the CLT provides a way of understanding the large-scale behaviour of such data from the small-scale individual parts.

In the world of insurance, CLT is at the foundation of risk diversification: in short, the idea that the aggregation of many not-completely known risks will result in a well-behaved average. To be valid, the CLT requires that the means and standard deviations of the distribution underlying each individual random variable converge to a stable value. However, as we saw above, this is not the case with power laws - and even when this is the case, convergence can be very slow. Difficulties with convergence in the CLT also arise for less pathological distributions such as the log-normal. Most insurers and reinsurers invest heavily in CAT models and analytics: little effort, however, is put into understanding the behaviour of the combination of CAT risks.

A practical consequence of this is that insurers, and especially reinsurers, are forever surprised by large events such as the Christchurch earthquakes or Thailand floods. The prevailing wisdom in the industry is that these surprises are a consequence of incomplete knowledge about nature. This, however, is only part of the truth. In contrast, while not every risk factor for car accidents can be quantified, every year their aggregated annual cost is unsurprisingly constant. To understand that CAT losses conform to a distinctly different statistical world altogether is another piece of the puzzle that the insurance industry has been slow to acknowledge.

References

Kousky, C. and Cooke, R. M. (2009) *The Unholy Trinity: Fat Tails, Tail Dependence and Micro-Correlations*, Resources for the Future 09-36-REV (http://www.rff.org/files/sharepoint/WorkImages/Download/RFF-DP-09-36-REV.pdf)

Mitzenmacher, M. (2003), A Brief History of Generative Models for Power Law and Lognormal distributions, *Internet Mathematics* 1(2) 226-251

Newman, M. E. J. (2006), Power laws, Pareto distributions and Zipf's law, arXiv:cond-mat/0412004, (http://arxiv.org/abs/cond-mat/0412004)

Auctions and the Government's Emissions Reduction Fund

by Rob van den Honert

Background

In searching for a sustainable policy to reduce the possible impacts of climate change, both sides of Federal politics support a reduction of Australia's carbon emissions to a level 5 per cent below 2000 levels by 2020. Labor and the Coalition both agree on the target, but disagree on the mechanism to employ to reach that target.

The Coalition's "Direct Action" climate change policy is centred on the establishment of a fund, the Emissions Reduction Fund (ERF). This is, in effect, a form of emissions trading scheme (ETS), in which businesses compete to win tenders (and are paid) to undertake emission reduction projects. The policy, introduced in a White Paper released in April 2014, has budgeted $2.55 billion to the ERF over the first four years.

The Coalition's policy replaces Labor's Carbon Tax, which taxed the country's biggest polluters, mostly coal-fired electricity generators (Figure 1). Polluters passed these increased taxes on to consumers. Tax-paying consumers were in turn compensated by the Government to help mitigate any price increases, such as increases in the cost of electricity.

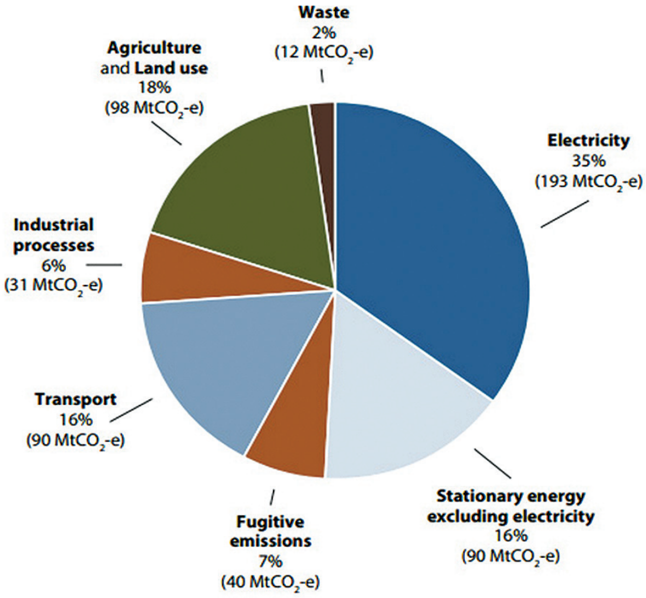

Source: Australian Government, *Australia's National Inventory Report 2012*, 2014.
Note: Figures are expressed using Kyoto Protocol accounting rules in terms of millions of tonnes of carbon dioxide equivalents ($MtCO_2$-e), using the global warming potentials published in the Intergovernmental Panel on Climate Change's Second Assessment Report.

Figure 1: Australia's major sources of carbon emissions

In a sense emissions trading is just the opposite of a carbon tax. Under an ETS, the amount of reduced carbon emissions required in any period is fixed by the government and the market then sets the price; under a carbon tax, the price of carbon emissions is fixed and polluters pay for the emissions they produce. This provides the polluters an incentive to reduce emissions.

Despite the Coalition Government's claims that the scheme will ensure a reduction of emissions for the best possible price, some critics claim that a carbon tax is cheaper and easier to implement, with lower administration and compliance costs. A fixed carbon price will give stability in the market and allow polluters to easily determine the viability of new and cleaner technology investments to assist in reducing their emissions. Some critics of the government's direct action approach argue that taxation is more direct and transparent than emissions trading, and affords less opportunity for gaming, speculation or corruption; moreover, taxation money moves from the polluters directly to the government (e.g. Hodgkinson & Johnston, 2015).

So how does the Government's Emissions Reduction Fund work?

The ERF operates as a reverse auction, in which businesses compete with each other to win a contract, and with it, Government monies. The Clean Energy Regulator runs auctions quarterly: a business nominates how much carbon pollution it is willing to reduce (in tonnes), together with a price per tonne. The Government, represented by the Clean Energy Regulator, awards contracts to the bidders who offer to reduce carbon at the cheapest price. The aim is to distribute the allocated budget in such a way to achieve maximum carbon reduction above a certain base level.

Amongst the set of bidders the maxi-min result is the desired outcome, i.e. the highest price of the bidders' lowest offers.

Mathematically the Government's distribution of the ERF budget can be formulated as a linear programming problem:

Let B = Quarterly ERF budget

X_i = 0 if project *i* is not selected

= 1 if project *i* is selected

Then the objective function should be

Maximise $\sum_{\text{All bids } i} [\text{Carbon reduction (tonnes)}_i * X_i]$

Subject to $\sum_{\text{All bids } i} [\text{Carbon reduction (tonnes)}_i * \text{Price/Tonne (\$/tonne)}_i * X_i] < B$

The constraint here is the upper limit, B, of the ERF's budget.

Other constraints, related to a lower bound on the number of tonnes of carbon reduction required to meet the Government's Renewable Energy Targets, an upper bound on the price it is prepared to pay for a tonne of carbon reduced or some geographical spread of projects, for example, could be added to this model.

Around 200 projects registered to participate in the first auction in April 2015, most of them farmers and waste management operators. Farmers might be first movers here if they view the ERF auction as an insurance policy against seasonal farming volatility, providing a stable form of income. Projects could include planting native trees, which store carbon, as well as agreeing to land management restrictions that might reduce livestock grazing. The ERF could be seen as compensation for farming losses.

This approach is not new in Government policy in Australia - the reverse auction method mirrors the existing National Water Market, which conducts water buybacks to increase base river flows. Similarly requests for tenders from outside contractors and consultants, commonplace in all government departments in Australia, are a form of reverse auction – a first-price sealed-bid reverse auction – bidders can only submit one bid each, and cannot see the bids of other participants so they cannot adjust their own bids accordingly.

Technically, a reverse auction is a type of auction in which the traditional roles of buyer and seller are reversed, with the primary objective to drive purchase prices downward - in this case, to reduce Government costs (Schoenherr & Mabert, 2007). In an ordinary auction (also known as a forward auction), buyers compete to obtain a good or service by offering increasingly higher prices: in a reverse auction the sellers (in this case the polluters) compete to obtain business from the buyer (the Government), and prices will typically decrease as the sellers undercut each other.

Information transparency, achieved through the process of all competing bids being revealed to all participating sellers in real time and allowing sellers to make multiple offers in response to competing offers, has been put forward as a mechanism that improves the chances of reaching the fair market value (Shalev and Asbjornsen, 2010). But information transparency is lacking in the Coalition's ERF - sellers have to make their bids without knowledge of other sellers' proposed carbon reductions and prices. This lack of information transparency reduces the dynamic nature of the bidding process, and this may reduce the likelihood of reaching the fair market value. Indeed, the formulation of the problem shows that the role of the auction is simply to gather information about a potential set of projects (carbon reductions and prices). If all bids were to be public then we surmise that bidders could strategise to achieve the best outcome for themselves, whilst still satisfying the Government's requirements for meeting their Renewable Energy Targets.

Alternative forms of auctions

In economic theory, an auction refers to any mechanism or set of trading rules for exchanging goods or services. The ERF is simply an auction, and auctions are the way that governments undertake a lot of business with suppliers. So this begs the question as to whether some other form of auction would better achieve the information transparency missing from the ERF.

The most common form of auction is the forward open ascending auction. This is seen most commonly in the sale of real estate in Australia. Competing potential buyers bid openly against one another generally with an auctioneer announcing (and encouraging) bids. Each bid in the sequence is required to be higher than the previous bid. The auction ends when no participant is willing to bid further, at which point the highest bidder pays their bid price. In this type of auction the current highest bid, and sequence of bids, is always publically available, maximising information transparency. But the objective here is to identify the maximum bid price, so this type of auction is not appropriate for the ERF.

The reverse auction in the ERF is similar in some respects to an open descending price auction, otherwise known as the Dutch auction. The Dutch auction is named for the Dutch tulip and cut flower auctions held daily in several centres in the Netherlands. At the Aalsmeer Flower Auction, near Amsterdam, the biggest commercial building in the world sees more than 21 million flowers from all over the world traded this way every day. Here the auction begins with a high asking price per stem for some quantity of similar flowers on display: the price is rapidly lowered until a buyer is willing to accept the price for some quantity of the flowers in the lot. If the first bidder does not purchase the entire lot, the price continues to lower until all of the flowers have been bid for.

In a single-lot Dutch auction only one bidder gets to place a bid - the others are all too late! But if there are multiple lots on auction, buyers are competing in real time against one another, and their bids (and hence strategies) are publically revealed. Modern technology has replaced the hustle and bustle of the market – bidding is now done online, and photographs of the lots of flowers are posted online, so bidders can be globally dispersed. This procedure has been going on in some form for more than a hundred years, and its survival to the present day attests to its robustness and efficiency. A key difference to the carbon market is that the flower growers want to maximise the price they will be paid, whilst the ERF wants to minimise the price they will have to pay for carbon. If all tender submissions for projects are made public then the ERF will reduce to a multi-lot Dutch auction if and only if the Government is obliged to accept the highest price bids - something that they are trying to avoid.

A modification of the Dutch auction is the reverse sealed-bid second-price auction, or reverse Vickrey auction. Here winning bidders (i.e. those with the lowest carbon price) are paid the second-lowest bid price rather than their own price. In the carbon pricing context where the ERF is searching to fund multiple projects, this would encourage downward price pressure on bidders, as there would be assurance that they would be paid more than their bid. On the other hand this mechanism may encourage collusion between project proponents to set prices.

Finally, from a purely economic point of view, a Walrasian auction might prove to be the fairest way of achieving a true market price for the ERF. This type of auction is a simultaneous auction in which buyers and sellers each calculate their own demand for a good at a range of possible prices, and each submits this to the "auctioneer". The price is then set so that the total demand across all buyers equals the total supply of the good. Thus a Walrasian auction perfectly matches supply and demand. In the context of the ERF all project proponents should submit a range of projects (i.e. tonnes of carbon reduced) and the price ($/tonne) they would accept for each of these projects, based on their own circumstances. This information will allow a market supply function to be determined. Simultaneously the Government should determine its demand function based on budgets and Renewable Energy Targets. The Walrasian auction result will be the equilibrium point, and signals the market price for carbon. Of course, whether this could be operationalized in practice remains to be seen.

Conclusion

A scan of the literature on auctions reveals that the scheme used to determine a market price for carbon reduction is probably good from the Government's perspective. The model presented here then allows the Government to meet their objectives in terms of emissions targets and budgets yet retains the opportunity to satisfy their own political policies. Only the market will determine whether or not the scheme will be successful in achieving its goals.

However, the choice of projects selected for funding through the ERF should be carefully considered. Simply paying farmers to lock up carbon in plantations of native trees in a bushfire-prone country like Australia may not be so smart. Furthermore it should be questioned whether the scheme is well targeted, given that Australia's major source of carbon emissions is the coal-fired electricity generation industry (Figure 1). These businesses should be further incentivised (or coerced) to contribute to the emissions reduction imperative.

References

Hodgkinson, D. & Johnston, R. (2015) Politics aside, a simple carbon tax makes more sense than a convoluted emissions trading scheme. The Conversation, 31 July 2015.

Schoenherr, T. & Mabert, V.A. (2007), Online reverse auctions: Common myths versus evolving reality. Business Horizons (Kelley School of Business, Indiana University) 50 (5), 373–384, doi:10.1016/j.bushor.2007.03.003.

Shalev, M.E. & Asbjornsen, S. (2010). Electronic reverse auctions and the public sector - factors of success. J. Public Procurement 10 (3), 428–452.

Christenson's Flower Auction (2015). Available at http://www.flowerauction.com.au/index.php?page=dutch. Downloaded 05 November 2015.